Minimax Optimal Procedures for Locally Private Estimation

John C. Duchi[†] Michael I. Jordan[*] Martin J. Wainwright[*]
jduchi@stanford.edu jordan@stat.berkeley.edu wainwrig@berkeley.edu

Stanford University[†] University of California, Berkeley[*]
Stanford, CA 94305 Berkeley, CA 94720

arXiv:1604.02390v1 [math.ST] 8 Apr 2016

Abstract

Working under a model of privacy in which data remains private even from the statistician, we study the tradeoff between privacy guarantees and the risk of the resulting statistical estimators. We develop private versions of classical information-theoretic bounds, in particular those due to Le Cam, Fano, and Assouad. These inequalities allow for a precise characterization of statistical rates under local privacy constraints and the development of provably (minimax) optimal estimation procedures. We provide a treatment of several canonical families of problems: mean estimation and median estimation, multinomial probability estimation, and nonparametric density estimation. For all of these families, we provide lower and upper bounds that match up to constant factors, and exhibit new (optimal) privacy-preserving mechanisms and computationally efficient estimators that achieve the bounds. Additionally, we present a variety of experimental results for estimation problems involving sensitive data, including salaries, censored blog posts and articles, and drug abuse; these experiments demonstrate the importance of deriving optimal procedures.

1 Introduction

A major challenge in statistical inference is that of characterizing and balancing statistical utility with the privacy of individuals from whom data is obtained [16, 17, 23]. Such a characterization requires a formal definition of privacy, and *differential privacy* has been put forth as one such candidate (e.g., see the papers [19, 8, 20, 28, 29] and references therein). In the database and cryptography literatures from which differential privacy arose, early research was mainly algorithmic in focus, and researchers have used differential privacy to evaluate privacy-retaining mechanisms for transporting, indexing, and querying data. More recent work aims to link differential privacy to statistical concerns [18, 50, 27, 45, 11, 43]; in particular, researchers have developed algorithms for private robust statistical estimators, point and histogram estimation, and principal components analysis. Much of this line of work is non-inferential in nature: as opposed to studying performance relative to an underlying population, they instead aim to approximate a class of statistics under privacy-respecting transformations for a fixed underlying data set. There has also been recent work within the context of classification problems and the "probably approximately correct" framework of statistical learning theory [e.g., 32, 7] that treats the data as random and aims to recover aspects of the underlying population.

In this paper, we take a fully inferential point of view on privacy, bringing differential privacy into contact with statistical decision theory. Our focus is on the fundamental limits of differentially-private estimation, and the identification of optimal mechanisms for enforcing a given level of privacy. By treating differential privacy as an abstract constraint on estimators, we obtain independence from specific estimation procedures and privacy-preserving mechanisms. Within this

framework, we derive both lower bounds and matching upper bounds on minimax risk. We obtain our lower bounds by integrating differential privacy into the classical paradigms for bounding minimax risk via the inequalities of Le Cam, Fano, and Assouad, while we obtain matching upper bounds by proposing and analyzing specific private procedures.

Differential privacy provides one formalization of the notion of "plausible deniability": no matter what public data is released, it is nearly equally as likely to have arisen from one underlying private sample as another. It is also possible to interpret differential privacy within a hypothesis testing framework [50], where the differential privacy parameter α controls the error rate in tests for the presence or absence of individual data points in a dataset (see Figure 3 for more details). Such guarantees against discovery, together with the treatment of issues of side information or adversarial strength that are problematic for other formalisms, have been used to make the case for differential privacy within the computer science literature; see, for example, the papers [22, 19, 5, 25]. In this paper, we study the setting of *local privacy*, in which providers do not even trust the statistician collecting the data. Although local privacy is a relatively stringent requirement, we view this setting as a natural step in identifying minimax risk bounds under privacy constraints. Indeed, local privacy is one of the oldest forms of privacy: its essential form dates back to Warner [49], who proposed it as a remedy for what he termed "evasive answer bias" in survey sampling.

Although differential privacy provides an elegant formalism for limiting disclosure and protecting against many forms of privacy breach, it is a stringent measure of privacy, and it is conceivably overly stringent for statistical practice. Indeed, Fienberg et al. [24] criticize the use of differential privacy in releasing contingency tables, arguing that known mechanisms for differentially private data release can give unacceptably poor performance. As a consequence, they advocate—in some cases—recourse to weaker privacy guarantees to maintain the utility and usability of released data. There are results that are more favorable for differential privacy; for example, Smith [45] shows that the non-local form of differential privacy [19] can be satisfied while yielding asymptotically optimal parametric rates of convergence for some point estimators. Resolving such differing perspectives requires investigation into whether particular methods have optimality properties that would allow a general criticism of the framework, and characterizing the trade-offs between privacy and statistical efficiency. Such are the goals of the current paper.

1.1 Our contributions

In this paper, we provide a formal framework for characterizing the tradeoff between statistical utility and local differential privacy. Basing our work on the classical minimax framework, our primary goals are to characterize how, for various types of estimation problems, the optimal rate of estimation varies as a function of the privacy level and other problem parameters. Within this framework, we develop a number of general techniques for deriving minimax bounds under local differential privacy constraints.

These bounds are useful in that they not only characterize the "statistical price of privacy," but they also allow us to compare different concrete procedures (or privacy mechanisms) for producing private data. Most importantly, our minimax theory can be used to identify which mechanisms are optimal, meaning that they preserve the maximum amount of statistical utility for a given privacy level. In practice, we find that these optimal mechanisms often differ from widely-accepted procedures from the privacy literature, and lead to better statistical performance while providing the same privacy guarantee. As one concrete example, Figure 1 provides an illustration of such gains in the context of estimating proportions of drug users based on privatized data. (See Section 6.2 for

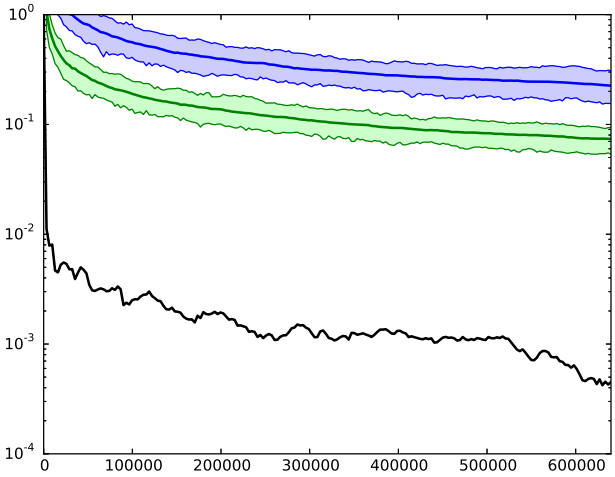

Figure 1. Estimating proportions of drug use, plotting the maximum error $\|\widehat{\theta} - \theta\|_\infty$ versus sample size. Top (blue) line: mean error of Laplace noise addition with 90% coverage intervals. Middle (green) line: mean error of minimax-optimal ℓ_∞ sampling strategy (28) with 90% coverage intervals. Bottom (black) line: mean error of non-private estimate.

a full description of the data set, and how these plots were produced.) The black curve shows the average ℓ_∞-error of a non-private estimator, based on access to the raw data; any estimator that operates on private data must necessarily have larger error than this gold standard. The upper two curves show the performance of two types of estimators that operate on privatized data: the blue curve is based on the standard mechanism of adding Laplace-distributed noise to the data, whereas the green curve is based on the optimal privacy mechanism identified by our theory. This optimal mechanism yields a roughly five-fold reduction in the mean-squared error, with the same computational requirements as the Laplacian-based procedure.

In this paper, we analyze the private minimax rates of estimation for several canonical problems: (a) mean estimation; (b) median estimation; (c) high-dimensional and sparse sequence estimation; (d) multinomial probability estimation; (e) conditional probability estimation; and (f) density estimation. To do so, we expand upon several canonical techniques for lower bounding minimax risk [52], establishing differentially private analogues of Le Cam's method in Section 3 and concomitant optimality guarantees for mean, median, and conditional probability estimators; Fano's method in Section 4, where we provide optimal procedures for high-dimensional estimation; and Assouad's method in Section 5, in which we investigate multinomial and density estimation. In accordance with our connections to statistical decision theory, we provide minimax rates for estimation of *population* quantities; by way of comparison, most prior work in the privacy literature focuses on accurate approximation of statistics in a conditional analysis in which the data are treated as fixed.

Notation: For distributions P and Q defined on a space $\mathcal{X}$, each absolutely continuous with respect to a measure μ (with corresponding densities p and q), the KL divergence between P and Q is

$$D_{\rm kl}\left(P\|Q\right) := \int_{\mathcal{X}} dP \log \frac{dP}{dQ} = \int_{\mathcal{X}} p \log \frac{p}{q} d\mu.$$

Letting $\sigma(\mathcal{X})$ denote an appropriate σ-field on $\mathcal{X}$, the total variation distance between P and Q is

$$\|P - Q\|_{\mathrm{TV}} := \sup_{S \in \sigma(\mathcal{X})} |P(S) - Q(S)| = \frac{1}{2} \int_{\mathcal{X}} |p(x) - q(x)| \, d\mu(x).$$

Given a pair of random variables (X, Y) with joint distribution $P_{X,Y}$, their mutual information is given by $I(X;Y) = D_{\mathrm{kl}}\left(P_{X,Y} \| P_X P_Y\right)$, where P_X and P_Y denote the marginal distributions. A random variable Y has the Laplace(α) distribution if its density is $p_Y(y) = \frac{\alpha}{2} \exp\left(-\alpha|y|\right)$. For matrices $A, B \in \mathbb{R}^{d \times d}$, the notation $A \preceq B$ means that $B - A$ is positive semidefinite. For sequences of real numbers $\{a_n\}$ and $\{b_n\}$, we use $a_n \lesssim b_n$ to mean there is a universal constant $C < \infty$ such that $a_n \leq C b_n$ for all n, and $a_n \asymp b_n$ to denote that $a_n \lesssim b_n$ and $b_n \lesssim a_n$. For a sequence of random variables X_n, we write $X_n \overset{d}{\rightsquigarrow} Y$ if X_n converges in distribution to Y.

2 Background and problem formulation

We begin by setting up the classical minimax framework, and then introducing the notion of an α-private minimax rate that we study in this paper.

2.1 Classical minimax framework

Let $\mathcal{P}$ denote a class of distributions on the sample space $\mathcal{X}$, and let $\theta(P) \in \Theta$ denote a function defined on $\mathcal{P}$. The space Θ in which the parameter $\theta(P)$ takes values depends on the underlying statistical model. For example, in the case of univariate mean estimation, Θ is a subset of the real line, whereas for a density estimation problem, Θ is some subset of the space of all possible densities over $\mathcal{X}$. Let ρ denote a semi-metric on the space Θ, which we use to measure the error of an estimator for the parameter θ, and let $\Phi : \mathbb{R}_+ \to \mathbb{R}_+$ be a non-decreasing function with $\Phi(0) = 0$ (for example, $\Phi(t) = t^2$).

In the non-private setting, the statistician is given direct access to i.i.d. observations $\{X_i\}_{i=1}^n$ drawn according to some distribution $P \in \mathcal{P}$. Based on the samples, the goal is to estimate the unknown parameter $\theta(P) \in \Theta$. We define an estimator $\widehat{\theta}$ as a measurable function $\widehat{\theta} : \mathcal{X}^n \to \Theta$, and we assess the quality of the estimate $\widehat{\theta}(X_1, \ldots, X_n)$ in terms of the risk

$$\mathbb{E}_P \left[\Phi\big(\rho(\widehat{\theta}(X_1, \ldots, X_n), \theta(P))\big) \right].$$

For instance, for a univariate mean problem with $\rho(\widehat{\theta}, \theta) = |\widehat{\theta} - \theta|$ and $\Phi(t) = t^2$, this risk is the mean-squared error. The risk assigns a nonnegative number to each pair $(\widehat{\theta}, \theta)$ of estimator and parameter.

The minimax risk is defined by the saddlepoint problem

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) := \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\Phi\big(\rho(\widehat{\theta}(X_1, \ldots, X_n), \theta(P))\big) \right], \tag{1}$$

where we take the supremum over distributions $P \in \mathcal{P}$ and the infimum over all estimators $\widehat{\theta}$. There is a substantial body of literature focused on techniques for upper- and lower-bounding the minimax risk for various classes of estimation problems. Our goal in this paper is to define and study a modified version of the minimax risk that accounts for privacy constraints.

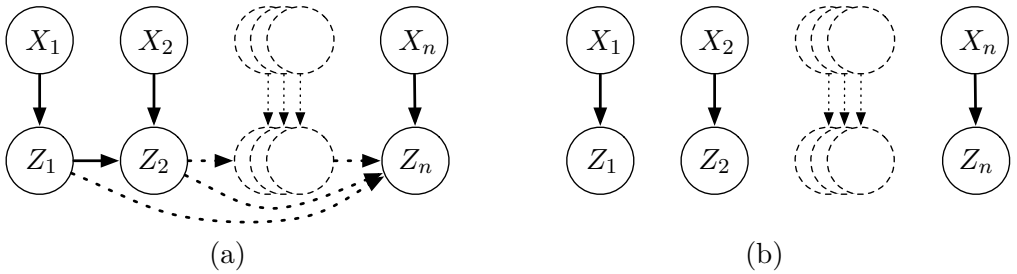

Figure 2. (a) Graphical model illustrating the conditional independence relationships between the private data $\{Z_i\}_{i=1}^n$ and raw variables $\{X_i\}_{i=1}^n$ in the interactive case. (b) Simpler graphical model illustrating the conditional independence structure in the non-interactive case.

2.2 Local differential privacy

Let us now define the notion of local differential privacy in a precise manner. The act of transforming data from the raw samples $\{X_i\}_{i=1}^n$ into a private set of samples $\{Z_i\}_{i=1}^n$ is modeled by a conditional distribution. We refer to this conditional distribution as either a *privacy mechanism* or a *channel distribution*, as it acts as a conduit from the original to the privatized data. In general, we allow the privacy mechanism to be *sequentially interactive*, meaning the channel has the conditional independence structure

$$\{X_i, Z_1, \ldots, Z_{i-1}\} \to Z_i \quad \text{and} \quad Z_i \perp X_j \mid \{X_i, Z_1, \ldots, Z_{i-1}\} \text{ for } j \neq i. \tag{2}$$

See panel (a) of Figure 2 for a graphical model that illustrates these conditional independence relationships. Given these independence relations (2), the full conditional distribution can be specified in terms of the conditionals $Q_i(Z_i \mid X_i = x_i, Z_{1:i-1} = z_{1:i-1})$. Thus, the full privacy mechanism can be specified by the collection $Q = \{Q_i\}_{i=1}^n$ of these conditional distributions. A special case is the *non-interactive* case, illustrated in panel (b) of Figure 2, in which each Z_i depends only on X_i. In this case, the conditional distributions take the simpler form $Q_i(Z_i \mid X_i = x_i)$.

The notion of local differential privacy involves placing some additional restrictions on the conditional distribution Q_i.

Definition 1. For a given privacy parameter $\alpha \geq 0$, the random variable Z_i is an *α-differentially locally private view* of X_i if for all $z_1, \ldots, z_{i-1}$ and $x, x' \in \mathcal{X}$ we have

$$\sup_{S \in \sigma(\mathcal{Z})} \frac{Q_i(Z_i \in S \mid X_i = x, Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1})}{Q_i(Z_i \in S \mid X_i = x', Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1})} \leq \exp(\alpha), \tag{3}$$

where $\sigma(\mathcal{Z})$ denotes an appropriate σ-field on $\mathcal{Z}$. We say that the privacy mechanism Q is α-differentially locally private (DLP) if each variable Z_i is an α-DLP view.

In the non-interactive case, the bound (3) reduces to

$$\sup_{S \in \sigma(\mathcal{Z})} \sup_{x, x' \in \mathcal{X}} \frac{Q(Z_i \in S \mid X_i = x)}{Q(Z_i \in S \mid X_i = x')} \leq \exp(\alpha). \tag{4}$$

The non-interactive version of local differential privacy dates back to the work of Warner [49]; see also Evfimievski et al. [22]. The more general interactive model was put forth by Dwork, McSherry,

Nissim, and Smith [19], and has been investigated in a number of works since then. In the context of our work on local privacy, relevant references include Beimel et al.'s [6] investigation of one-dimensional Bernoulli probability estimation under the model (3), and Kairouz et al.'s [31] study of channel constructions that maximize information-theoretic measures of information content for various domains $\mathcal{X}$.

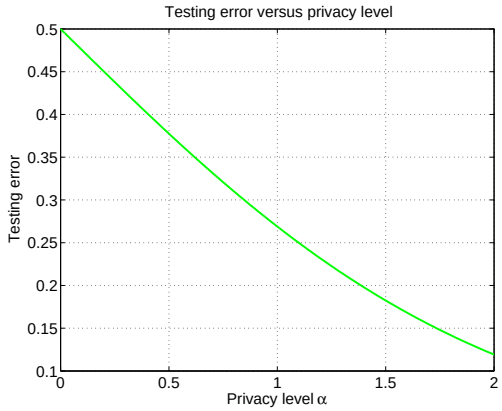

Figure 3. For any α-DLP view Z of X, the probability of error P_{err} in distinguishing between the two hypotheses $\{X = x\}$ and $\{X = x'\}$ is lower bounded as $P_{\mathrm{err}} \geq \frac{1}{1+e^{\alpha}}$. Thus, for $\alpha \in [0, \frac{1}{4}]$, we are guaranteed that $P_{\mathrm{err}} \geq 0.43$.

As discussed by Wasserman and Zhou [50], one intuitive interpretation of differential privacy is in terms of disclosure risk. More concretely, suppose that given an α-private view Z of the random variable X, our goal is to distinguish between the two hypotheses $\{X = x\}$ versus $\{X = x'\}$, where $x, x' \in \mathcal{X}$ are two distinct possible values in the sample space. A calculation shows that the best possible probability of error of any hypothesis test, with equal weights for each hypothesis, satisfies

$$P_{\mathrm{err}} := \frac{1}{2} \cdot \inf_{\psi} \left\{ \mathbb{P}(\psi(Z) \neq x \mid X = x) + \mathbb{P}(\psi(Z) \neq x' \mid X = x') \right\} \geq \frac{1}{1 + e^{\alpha}}.$$

Consequently, small values of α ensure that the performance of any test is close to random guessing (see Figure 3). We relate this in passing to Warner's classical randomized response mechanism [49] in a simple scenario with $X \in \{0, 1\}$, where we set $Z = X$ with probability $q_{\alpha} = \frac{e^{\alpha}}{1+e^{\alpha}}$, and $Z = 1 - X$ otherwise. Then $Q(Z = z \mid X = x)/Q(Z = z \mid X = x') \in [e^{-\alpha}, e^{\alpha}]$, and the disclosure risk P_{err} is precisely $1/(1 + e^{\alpha})$.

2.3 α-private minimax risks

Given our definition of local differential privacy (LDP), we are now equipped to describe the notion of an α-LDP minimax risk. For a given privacy level $\alpha > 0$, let $\mathcal{Q}_{\alpha}$ denote the set of all conditional distributions have the α-LDP property (3). For a given raw sample $\{X_i\}_{i=1}^n$, any distribution $Q \in \mathcal{Q}_{\alpha}$ produces a set of private observations $\{Z_i\}_{i=1}^n$, and we now restrict our attention to estimators $\widehat{\theta} = \widehat{\theta}(Z_1, \ldots, Z_n)$ that depend purely on this private sample. Doing so yields the modified minimax risk

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho; Q) := \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_{P,Q} \left[\Phi\big(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P))\big) \right], \tag{5}$$

where our notation now reflects its dependence on the choice of privacy mechanism Q. By definition, any choice of $Q \in \mathcal{Q}_\alpha$ guarantees that the data $\{Z_i\}_{i=1}^n$ are α-locally differentially private, so that it is natural to seek the mechanism(s) in $\mathcal{Q}_\alpha$ that lead to the smallest values of the minimax risk (5). This minimization problem leads to the central object of study for this paper, a functional which characterizes the optimal rate of estimation in terms of the privacy parameter α.

Definition 2. Given a family of distributions $\theta(\mathcal{P})$ and a privacy parameter $\alpha > 0$, the *α-private minimax risk* in the metric ρ is

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, \alpha) := \inf_{Q \in \mathcal{Q}_\alpha} \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_{P,Q} \left[\Phi(\rho(\widehat{\theta}(Z_1, \ldots, Z_n), \theta(P))) \right]. \tag{6}$$

Note that as $\alpha \to +\infty$, the constraint of membership in $\mathcal{Q}_\alpha$ becomes vacuous, so that the α-private minimax risk reduces to the classical minimax risk (1). Of primary interest in this paper are settings in which $\alpha \in (0, 1]$, corresponding to reasonable levels of the disclosure risk, as illustrated in Figure 3.

3 Bounds on pairwise divergences: Le Cam's bound and variants

Perhaps the oldest approach to bounding the classical minimax risk (1) is via Le Cam's method [35]. Beginning with this technique, we develop a private analogue of the Le Cam bound, and we show how it can be used to derive sharp lower bounds on the α-private minimax risk for one-dimensional mean and median estimation problems, as well as one-dimensional conditional probability or other association estimates. We also provide new optimal procedures for each of these settings.

3.1 A private version of Le Cam's bound

The classical (non-private) version of Le Cam's method bounds the classical minimax risk (1) in terms of a two-point hypothesis testing problem [35, 52, 47]. For any distribution P, we use P^n to denote the product distribution correpsonding to a collection of n i.i.d. samples. Let us say that a pair of distributions $\{P_1, P_2\}$ is 2δ-separated if $\rho(\theta(P_1), \theta(P_2)) \geq 2\delta$. With this terminology, a simple version of Le Cam's lemma asserts that, for any 2δ-separated pair of distributions, the classical minimax risk is lower bounded as

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \stackrel{(i)}{\geq} \frac{\Phi(\delta)}{2} \left\{ 1 - \|P_1^n - P_2^n\|_{\mathrm{TV}} \right\} \stackrel{(ii)}{\geq} \frac{\Phi(\delta)}{2} \left\{ 1 - \frac{1}{\sqrt{2}} \sqrt{n D_{\mathrm{kl}}\left(P_1 \| P_2\right)} \right\}. \tag{7}$$

Given the first lower bound (i), the second bound (ii) follows as a consequence of the Pinsker bound on the total variation norm in terms of the KL divergence

$$\|P_1^n - P_2^n\|_{\mathrm{TV}}^2 \leq \frac{1}{2} D_{\mathrm{kl}}\left(P_1^n \| P_2^n\right),$$

along with the fact $D_{\mathrm{kl}}\left(P_1^n \| P_2^n\right) = n D_{\mathrm{kl}}\left(P_1 \| P_2\right)$, which follows because the distributions P_1^n and P_2^n are product distributions (i.e., we have $X_i \stackrel{\mathrm{iid}}{\sim} P$).

Let us now state a version of Le Cam's lemma that applies to the α-locally private setting in which the estimator $\widehat{\theta}$ depends only on the private variables $(Z_1, \ldots, Z_n)$, and our goal is to lower bound the α-private minimax risk (6).

Proposition 1 (Private form of Le Cam bound)**.** *Suppose that we are given n i.i.d. observations from an α-locally differential private channel for some $\alpha \in [0, \frac{23}{35}]$. Then for any 2δ-separated pair of distributions (P_1, P_2), the α-private minimax risk has lower bound*

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, \alpha) \geq \frac{\Phi(\delta)}{2} \left\{1 - \sqrt{4\alpha^2 n \, \|P_1 - P_2\|_{\mathrm{TV}}^2}\right\} \; \geq \; \frac{\Phi(\delta)}{2} \left\{1 - \sqrt{2\alpha^2 n D_{\mathrm{kl}}\left(P_1 \| P_2\right)}\right\}. \quad (8)$$

Remarks: A comparison with the original Le Cam bound (7) shows that for $\alpha \in [0, \frac{23}{35}]$, the effect of α-local differential privacy is to reduce the *effective sample size* from n to at most $4\alpha^2 n$. Proposition 1 is a corollary of more general results, which we describe in Section 3.3, that quantify the contraction in KL divergence that arises from passing the data through an α-private channel.

3.2 Some applications of the private Le Cam bound

We now turn to some applications of the α-private version of Le Cam's inequality from Proposition 1. First, we study the problem of one-dimensional mean estimation. In addition to demonstrating how the minimax rate changes as a function of α, we also reveal some interesting (and perhaps disturbing) effects of enforcing α-local differential privacy: the effective sample size may be even polynomially smaller than $\alpha^2 n$. Our second example studies median estimation, which—as a more robust quantity than the mean—allows us to always achieve parametric convergence rates with an effective sample size reduction of n to $\alpha^2 n$. Our third example investigates conditional probability estimation, which exhibits a more nuanced dependence on privacy than the preceding estimates. We state each of our bounds assuming $\alpha \in [0, 1]$; the bounds hold (with different numerical constants) whenever $\alpha \in [0, C]$ for some universal constant C.

3.2.1 One-dimensional mean estimation

Given a real number $k > 1$, consider the family

$$\mathcal{P}_k := \big\{\text{distributions } P \text{ such that } \mathbb{E}_P[X] \in [-1, 1] \text{ and } \mathbb{E}_P[|X|^k] \leq 1\big\}.$$

Note that the parameter k controls the tail behavior of the random variable X, with larger values of k imposing more severe constraints. As $k \uparrow +\infty$, the random variable converges to one that is supported almost surely on the interval $[-1, 1]$. Suppose that our goal is to estimate the mean $\theta(P) = \mathbb{E}_P[X]$, and that we adopt the squared ℓ_2-error to measure the quality of an estimator. The classical minimax risk for this problem scales as $n^{-\min\{1, 2-\frac{2}{k}\}}$ for all values of $k \geq 1$. Our goal here is to understand how the α-private minimax risk (6),

$$\mathfrak{M}_n(\theta(\mathcal{P}_k), (\cdot)^2, \alpha) := \inf_{Q \in \mathcal{Q}_\alpha} \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}_k} \mathbb{E}\left[\left(\widehat{\theta}(Z_1, \ldots, Z_n) - \theta(P)\right)^2\right],$$

differs from the classical minimax risk.

Corollary 1. *There exist universal constants $0 < c_\ell \leq c_u \leq 9$ such that for all $k > 1$ and $\alpha \in [0, 1]$, the α-minimax risk $\mathfrak{M}_n(\theta(\mathcal{P}_k, (\cdot)^2, \alpha)$ is sandwiched as*

$$c_\ell \min\left\{1, \left(n\alpha^2\right)^{-\frac{k-1}{k}}\right\} \; \leq \; \mathfrak{M}_n(\theta(\mathcal{P}_k), (\cdot)^2, \alpha) \; \leq \; c_u \min\left\{1, \left(n\alpha^2\right)^{-\frac{k-1}{k}}\right\}. \quad (9)$$

We prove the lower bound using the α-private version (8) of Le Cam's inequality from Proposition 1; see Appendix B.1 for the details.

In order to understand the bound (9), it is worthwhile considering some special cases, beginning with the usual setting of random variables with finite variance ($k = 2$). In the non-private setting in which the original sample $(X_1, \ldots, X_n)$ is observed, the sample mean $\widehat{\theta} = \frac{1}{n}\sum_{i=1}^n X_i$ has mean-squared error at most $1/n$. When we require α-local differential privacy, Corollary 1 shows that the minimax rate worsens to $1/\sqrt{n\alpha^2}$. More generally, for any $k > 1$, the minimax rate scales as $\mathfrak{M}_n(\theta(\mathcal{P}_k), (\cdot)^2, \alpha) \asymp (n\alpha^2)^{-\frac{k-1}{k}}$. As $k \uparrow \infty$, the moment condition $\mathbb{E}[|X|^k] \le 1$ becomes equivalent to the boundedness constraint $|X| \le 1$ a.s., and we obtain the more standard parametric rate $(n\alpha^2)^{-1}$, where there is no reduction in the exponent.

The upper bound is achieved by a variety of privacy mechanisms and estimators. One of them is the following variant of the Laplace mechanism:

- Letting $[x]_T = \max\{-T, \min\{x, T\}\}$ denote the projection of x to the interval $[-T, T]$, output the private samples

$$Z_i = [X_i]_T + W_i, \quad \text{where} \quad W_i \stackrel{\text{iid}}{\sim} \text{Laplace}(\alpha/(2T)) \quad \text{and } T = (n\alpha^2)^{\frac{1}{2k}}. \tag{10}$$

- Compute the sample mean $\widehat{\theta}_n := \frac{1}{n}\sum_{i=1}^n Z_i$ of the privatized data.

We present the analysis of this estimator in Appendix B.1. It is one simple case in which the widely-used Laplacian mechanism is an optimal mechanism; later examples show that this is *not* always the case.

Summarizing our results thus far, Corollary 1 helps to demarcate situations in which local differential privacy may or may not be acceptable for location estimation problems. In particular, for bounded domains—where we may take $k \uparrow \infty$—local differential privacy may be quite reasonable. However, in situations in which the sample takes values in an unbounded space, local differential privacy imposes more severe constraints.

3.2.2 One-dimensional median estimation

Instead of attempting to privately estimate the mean—an inherently non-robust quantity—we may also consider median estimation problems. Median estimation for general distributions is impossible even in non-private settings,[1] so we focus on the median as an M-estimator. Recalling that the minimizer(s) of $\mathbb{E}[|X - \theta|]$ are the median(s) of X, we consider the gap between mean absolute error of our estimator and that of the true median,

$$\mathbb{E}[R(\widehat{\theta})] - \inf_{\theta \in \mathbb{R}} R(\theta), \quad \text{where} \quad R(\theta) := \mathbb{E}[|X - \theta|].$$

We first give a proposition characterizing the minimax rate for this problem by applying Proposition 1. Let $\theta(P) = \text{med}(P)$ denote the median of the distribution P, and for radii $r > 0$, we consider the family of distributions supported on $\mathbb{R}$ defined by

$$\mathcal{P}_r := \{\text{distributions } P \text{ such that } |\,\text{med}(P)| \le r, \ \mathbb{E}_P[|X|] < \infty\}.$$

[1] That is, the minimax error never converges to zero: consider estimating the median of the two distributions P_0 and P_1, each supported on $\{-r, r\}$, where $P_0(r) = \frac{1+\delta}{2}$ and $P_1(r) = \frac{1-\delta}{2}$, then take $\delta \downarrow 0$ as the sample size increases.

In this case, we consider the slight variant of the minimax rate (6) defined by the risk gap

$$\mathfrak{M}_n(\theta(\mathcal{P}_r), R, \alpha) := \inf_{Q \in \mathcal{Q}_\alpha} \inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}_r} \mathbb{E}_{P,Q} \left[R(\widehat{\theta}(Z_1, \ldots, Z_n)) - R(\theta(P)) \right].$$

We then have the following.

Corollary 2. *For the median estimation problem, there are universal constants $\frac{1}{20} \le c_\ell \le c_u < 6$ such that for all $r \ge 0$ and $\alpha \in [0, 1]$, the minimax error satisfies*

$$c_\ell \cdot r \min\left\{1, (n\alpha^2)^{-\frac{1}{2}}\right\} \le \mathfrak{M}_n(\theta(\mathcal{P}_r), R, \alpha) \le c_u \cdot r \min\left\{1, (n\alpha^2)^{-\frac{1}{2}}\right\}.$$

We present the proof of the lower bound in Corollary 2 to Section B.2, focusing our attention here on a minimax optimal sequential procedure based on stochastic gradient descent (SGD).

To describe our SGD procedure, let $\theta_0 \in [-r, r]$ be arbitrary, and W_i be an i.i.d. $\{-1, +1\}$ Bernoulli sequence with $\mathbb{P}(W_i = 1) = \frac{e^\alpha}{e^\alpha + 1}$, and let $X_i \stackrel{\text{iid}}{\sim} P$ be the observations of the distribution P whose median we wish to estimate (and which must be made private). We iterate according to the projected stochastic gradient descent procedure

$$\theta_{i+1} = \left[\theta_i - \eta_i Z_i\right]_r, \quad \text{where} \quad Z_i = \frac{e^\alpha + 1}{e^\alpha - 1} \cdot W_i \cdot \operatorname{sign}(\theta_i - X_i), \tag{11}$$

where as in expression (10), $[x]_r = \max\{-r, \min\{x, r\}\}$ is the projection onto the set $[-r, r]$, and the sequence $\eta_i > 0$ are non-increasing stepsizes. By inspection we see that Z_i is differentially private for X_i, and we have the conditional unbiasedness $\mathbb{E}[Z_i \mid X_i, \theta_i] = \operatorname{sign}(\theta_i - X_i) \in \partial_\theta |\theta_i - X_i|$, where ∂ denotes the subdifferential operator. Standard results on stochastic gradient descent methods [40] imply that for $\widehat{\theta}_n = \frac{1}{n} \sum_{i=1}^n \theta_i$, we have

$$\mathbb{E}[R(\widehat{\theta}_n)] - \inf_{\theta \in [-r,r]} R(\theta) \le \frac{1}{n} \left[\frac{r^2}{\eta_n} + \frac{1}{2} \sum_{i=1}^n \eta_i \left(\frac{e^\alpha + 1}{e^\alpha - 1} \right)^2 \right].$$

Under the assumption that $\alpha \le 1$, we take $\eta_i = \alpha \cdot r / \sqrt{i}$, which immediately implies the upper bound $\mathbb{E}[R(\widehat{\theta}_n)] - R(\operatorname{med}(P)) \le \frac{6r}{\sqrt{n\alpha^2}}$.

We make two remarks on the procedure (11). First, it is essentially a sequential variant of Warner's 1965 randomized response [49], a procedure whose variants are often optimal, as we show in the sequel. Secondly, while at first blush it is not clear that the additional complexity of stochastic gradient descent is warranted, we provide experiments comparing the SGD procedure with more naive estimators in Section 6.1.2 on a salary estimation task. These experiments corroborate the improved performance of our minimax optimal strategy.

3.2.3 Conditional probability estimation

An important problem in predictive tasks—such as as genome-wide association studies [e.g., 12] or in generative modeling for classification [e.g., 30, Chapter 6.6]—is the estimation of the conditional probability of a signal Y given single covariates. In particular, we wish to estimate the conditional probability model $P(Y = y \mid X = x)$ for binary $x \in \{0, 1\}$ and $y \in \{-1, 1\}$. The two-parameter logistic model

$$P(Y = y \mid X = x) = \frac{1}{1 + \exp(-y(\theta_0 + \theta_1 x))},$$

with parameters θ_0 and θ_1 captures this setting. Given an estimated conditional model $\widehat{P}$ and associated estimate $\widehat{\theta}$, the standard representation of θ as the log-odds implies that

$$c\|\widehat{\theta} - \theta\|_2 \le \max_{x\in\{0,1\}, y\in\{-1,1\}} \left| \log \frac{P(Y = y \mid X = x)}{\widehat{P}(Y = y \mid X = x)} \right| \le C\|\widehat{\theta} - \theta\|_2,$$

for positive numerical constants $c \le C$. That is, the conditional probability estimate has uniformly small error if and only if $\|\widehat{\theta} - \theta\|_2$ is small.

We consider minimax estimation for a family of problems with slightly restricted log-odds. Specifically, we consider the two-parameter family of conditional probability distributions indexed by the parameter bound r and probability $\mu \in \left(0, \frac{1}{2}\right]$ of seeing a non-zero x:

$$\mathcal{P}_{r,\mu} := \left\{ P \text{ s.t. } P(Y = y \mid X = x) = \frac{1}{1 + e^{-y(\theta_0 + \theta_1 x)}}, \;\; \|\theta\|_2 \le r, \;\; \text{and } P(X = 1) \in [\mu, 1 - \mu] \right\}.$$

We then have the following minimax guarantees.

Corollary 3. *For the conditional probability estimation problem (logistic parameter estimation problem), for all $\alpha \in [0, 1]$ the minimax estimation rate satisfies*

$$\mathfrak{M}_n(\theta(\mathcal{P}_{r,\mu}), \|\cdot\|_2, \alpha) \ge \frac{1}{8} \min \left\{ r, \log \left(1 + \frac{e^r}{4\sqrt{n\alpha^2 \min\{\mu^2, (1-\mu)^2\}}} \right) \right\}. \tag{12a}$$

Conversely, let $\alpha \ge 0$ and $q_\alpha = \frac{e^{\alpha/2}}{1+e^{\alpha/2}}$. There exists an α-locally-differentially private estimator $\widehat{\theta}_n$ such that for any logistic distribution P_θ,

$$n\|\widehat{\theta}_n - \theta\|_2^2 \overset{d}{\rightsquigarrow} C_{\alpha,\theta}^2 \cdot \chi_1^2, \tag{12b}$$

where χ_1^2 is a χ^2-random variable with 1 degree of freedom, and the constant $C_{\alpha,\theta}$ satisfies

$$C_{\alpha,\theta}^2 \asymp \max \left\{ \frac{e^{|\theta_0|}}{1-\mu}, \; \frac{e^{|\theta_1|}}{\mu}, \; \frac{q_\alpha(1-q_\alpha)}{(2q_\alpha - 1)^2} \cdot \frac{e^{2|\theta_0|}}{(1-\mu)^2}, \; \frac{q_\alpha(1-q_\alpha)}{(2q_\alpha - 1)^2} \cdot \frac{e^{2|\theta_1|}}{\mu^2} \right\}.$$

We prove the lower bound using Proposition 1; see Appendix B.3 for a proof of both equations (12).

Unpacking the lower bound (12a), we see that asymptotically as $n\alpha^2 \to \infty$, the lower bound on minimax risk under mean-squared error scales as $e^{2r}/(n\mu^2\alpha^2)$, which matches the constant $C_{\alpha,\theta}$ in the upper bound (12b) for small α, as $q_\alpha(1-q_\alpha)/(2q_\alpha - 1)^2 \approx 4/\alpha^2$. The dependence of the bound on the probability $P(X = x) = \mu$ and the exponentiated radius e^r is subtle; in the non-private case—as demonstrated by the bound (12b) with $\alpha = +\infty$ and $q_\alpha = 1$—the scaling in n depends linearly on μ as, intuitively, the observations "useful" for inference about θ_1 occur on the μ fraction of observations when $X = 1$ and those "useful" for inference about θ_0 occur on the $1 - \mu$ fraction of observations when $X = 0$. In the private case, then, the effective sample size decreases by a factor $\min\{\mu, 1-\mu\}\alpha^2/e^r < 1$ for this problem. In contrast to the simpler linear estimators of the previous sections, which exhibited a degradation in effective sample size of $n \mapsto n\alpha^2$, we see that conditional estimation adds additional complexity.

For the logistic estimation problem, one minimax optimal estimator is based on the plug-in principle. Our procedure (similar to that of Lei [36]) privately estimates the joint distribution of $(X, Y) \in \{0, 1\} \times \{-1, 1\}$, then, given the estimate $\widehat{P}_n$ based on n private observations, chooses

$$\widehat{\theta}_n \in \operatorname*{argmin}_{\theta : \|\theta\|_2 \le r} \left\{ \mathbb{E}_{\widehat{P}_n} \left[\log(1 + \exp(-Y(\theta_0 + \theta_1 X))) \right] \right\}. \tag{13}$$

We describe our construction of $\widehat{P}_n$ briefly here, as we use the private multimonial estimation strategy described in Section 5.2.1. We treat the observations $(X_i, Y_i) \in \{0, 1\} \times \{-1, 1\}$ as draws from a multinomial distribution on the standard basis vectors $\{e_1, \ldots, e_4\}$ by defining

$$U(x, y) = \{e_1 \text{ if } y = 1, x = 0, \quad e_2 \text{ if } y = -1, x = 0, \quad e_3 \text{ if } y = 1, x = 1, \quad e_4 \text{ if } y = -1, x = 1\}.$$

Then $U(X, Y)$ is multinomial with parameter $p = \mathbb{E}[U]$. Our optimal locally private multinomial estimator $\widehat{P}_n$ (see Section 5.2.1 to come) is sufficient as a plug-in estimator (13) to attain the minimax optimal rate (12b).

3.3 Pairwise upper bounds on Kullback-Leibler divergences

As mentioned previously, the private form of Le Cam's bound (Proposition 1) is a corollary of more general results on the contractive effects of privacy on pairs of distributions, which we now state. Given a pair of distributions (P_1, P_2) defined on a common space $\mathcal{X}$, any conditional distribution Q transforms such a pair of distributions into a new pair (M_1, M_2) via the marginalization operation $M_j(S) = \int_{\mathcal{X}} Q(S \mid x) dP_j(x)$ for $j = 1, 2$. Intuitively, when the conditional distribution Q satisfies the constraint of α-local differential privacy, the two output distributions (M_1, M_2) should be closer together. The following theorem makes this intuition precise:

Theorem 1. *For any $\alpha \ge 0$, let Q be a conditional distribution that guarantees α-differential privacy. Then, for any pair of distributions P_1 and P_2, the induced marginals M_1 and M_2 satisfy the bound*

$$D_{\rm kl}\left(M_1 \| M_2\right) + D_{\rm kl}\left(M_2 \| M_1\right) \le \min\{4, e^{2\alpha}\}(e^\alpha - 1)^2 \left\|P_1 - P_2\right\|_{\rm TV}^2. \tag{14}$$

Remarks: Theorem 1 is a type of *strong data processing* inequality [2], providing a quantitative relationship between the divergence $\|P_1 - P_2\|_{\rm TV}$ and the KL-divergence $D_{\rm kl}\left(M_1 \| M_2\right)$ that arises after applying the channel Q. The result of Theorem 1 is similar to a result due to Dwork, Rothblum, and Vadhan [20, Lemma III.2], who show that $D_{\rm kl}\left(Q(\cdot \mid x) \| Q(\cdot \mid x')\right) \le \alpha(e^\alpha - 1)$ for any $x, x' \in \mathcal{X}$, which implies $D_{\rm kl}\left(M_1 \| M_2\right) \le \alpha(e^\alpha - 1)$ by convexity. This upper bound is weaker than Theorem 1 since it lacks the term $\|P_1 - P_2\|_{\rm TV}^2$. This total variation term is essential to our minimax lower bounds: more than providing a bound on KL divergence, Theorem 1 shows that differential privacy acts as a contraction on the space of probability measures. This contractivity holds in a strong sense: indeed, the bound (14) shows that even if we start with a pair of distributions P_1 and P_2 whose KL divergence is infinite, the induced marginals M_1 and M_2 always have finite KL divergence. We provide the proof of Theorem 1 in Appendix A.

Let us now develop a corollary of Theorem 1 that has a number of useful consequences, among them the private form of Le Cam's method from Proposition 1. Suppose that we are given an

indexed family of distributions $\{P_\nu, \nu \in \mathcal{V}\}$. Let V denote a random variable that is uniformly distributed over the finite index set $\mathcal{V}$. Conditionally on $V = \nu$, suppose we sample a random vector $(X_1, \ldots, X_n)$ according to a product measure of the form $P_\nu(x_1, \ldots, x_n) := \prod_{i=1}^n P_{\nu(i)}(x_i)$, where $(\nu(1), \ldots, \nu(n))$ denotes some sequence of indices. Now suppose that we draw an α-locally private sample $(Z_1, \ldots, Z_n)$ according to the channel $Q(\cdot \mid X_{1:n})$. Conditioned on $V = \nu$, the private sample is distributed according to the measure M_ν^n given by

$$M_\nu^n(S) := \int Q^n(S \mid x_1, \ldots, x_n) dP_\nu(x_1, \ldots, x_n) \quad \text{for } S \in \sigma(\mathcal{Z}^n). \tag{15}$$

Since we allow interactive protocols, the distribution M_ν^n need not be a product distribution in general. Nonetheless, in this setup we have the following tensorization inequality:

Corollary 4. *For any α-locally differentially private (3) conditional distribution Q and any paired sequences of distributions $\{P_{\nu(i)}\}$ and $\{P_{\nu'(i)}\}$, we have*

$$D_{\rm kl}\left(M_\nu^n \| M_{\nu'}^n\right) \leq 4(e^\alpha - 1)^2 \sum_{i=1}^n \left\| P_{\nu(i)} - P_{\nu'(i)} \right\|_{\rm TV}^2. \tag{16}$$

See Appendix A.2 for the proof, which requires a few intermediate steps to obtain the additive inequality. One consequence of Corollary 4 is the private form of Le Cam's bound from Proposition 1. Given the index set $\mathcal{V} = \{1, 2\}$, consider two paired sequences of distributions of the form $\{P_1, \ldots, P_1\}$ and $\{P_2, \ldots, P_2\}$. With this choice, we have

$$\|M_1^n - M_2^n\|_{\rm TV}^2 \overset{(i)}{\leq} \frac{1}{2} D_{\rm kl}\left(M_1^n \| M_2^n\right) \overset{(ii)}{\leq} 2(e^\alpha - 1)^2 n \left\| P_1 - P_2 \right\|_{\rm TV},$$

where step (i) follows from Pinsker's inequality, and step (ii) follows from the tensorization inequality (16) and the i.i.d. nature of the product distributions P_1^n and P_2^n. Noting that $(e^\alpha - 1)^2 \leq 2\alpha^2$ for $\alpha \leq \frac{23}{35}$ and applying the classical Le Cam bound (7) gives Proposition 1.

In addition, inequality (16) can be used to derive a bound on the mutual information. Bounds of this type are useful in applications of Fano's method, to be discussed at more length in the following section. In particular, if we define the mixture distribution $\overline{M}^n = \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} M_\nu^n$, then by the definition of mutual information, we have

$$\begin{aligned} I(Z_1, \ldots, Z_n; V) \; &= \; \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} D_{\rm kl}\left(M_\nu^n \| \overline{M}^n\right) \leq \frac{1}{|\mathcal{V}|^2} \sum_{\nu, \nu'} D_{\rm kl}\left(M_\nu^n \| M_{\nu'}^n\right) \\ &\leq 4(e^\alpha - 1)^2 \sum_{i=1}^n \frac{1}{|\mathcal{V}|^2} \sum_{\nu, \nu' \in \mathcal{V}} \left\| P_{\nu(i)} - P_{\nu'(i)} \right\|_{\rm TV}^2, \end{aligned} \tag{17}$$

the first inequality following from the joint convexity of the KL divergence and the final inequality from Corollary 4.

Remarks: Mutual information bounds under local privacy have appeared previously. McGregor et al. [38] study relationships between communication complexity and differential privacy, showing that differentially private schemes allow low communication. They provide a result [38, Prop. 7] guaranteeing $I(X_{1:n}; Z_{1:n}) \leq 3\alpha n$; they strengthen this bound to $I(X_{1:n}; Z_{1:n}) \leq (3/2)\alpha^2 n$ when

the X_i are i.i.d. uniform Bernoulli variables. Since the total variation distance is at most 1, our result also implies this scaling (for arbitrary X_i); however, our result is stronger since it involves the total variation terms $\|P_{\nu(i)} - P_{\nu'(i)}\|_{\mathrm{TV}}$. These TV terms are an essential part of obtaining the sharp minimax results that are our focus. In addition, Corollary 4 allows for *any* sequentially interactive channel Q; each Z_i may depend on the private answers $Z_{1:i-1}$ of other data providers.

4 Bounds on private mutual information: Fano's method

We now turn to a set of techniques for bounding the private minimax risk (6) based on Fano's inequality from information theory. We begin by describing how Fano's inequality is used in classical minimax theory, then presenting some of its extensions to the private setting.

Recall that our goal is to lower bound the minimax risk associated with estimating some parameter $\theta(P)$ in a given metric ρ. Given a finite set $\mathcal{V}$, a family of distributions $\{P_\nu, \nu \in \mathcal{V}\}$ is said to be 2δ-separated in the metric ρ if $\rho(\theta(P_\nu), \theta(P_{\nu'})) \geq 2\delta$ for all distinct pairs $\nu, \nu' \in \mathcal{V}$. Given any such 2δ-separated set, the classical form of Fano's inequality [cf. 52] asserts that the minimax risk (1) has lower bound

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \geq \frac{\Phi(\delta)}{2} \left\{1 - \frac{I(V; X_1^n) + \log 2}{\log |\mathcal{V}|}\right\}.$$

Here $I(V; X_1^n)$ denotes the mutual information between a random variable V uniformly distributed over the set $\mathcal{V}$ and a random vector $X_1^n = (X_1, \ldots, X_n)$ drawn from the mixture distribution

$$\overline{P} := \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} P_\nu^n, \tag{18}$$

so that $I(V; X_1^n) = \frac{1}{|\mathcal{V}|} \sum_\nu D_{\mathrm{kl}}\left(P_\nu^n \| \overline{P}\right)$; equivalently, the random variables are drawn $X_i \stackrel{\mathrm{iid}}{\sim} P_\nu$ conditional on $V = \nu$. In the cases we consider, it is sometimes convenient to use a slight generalization of the classical Fano method by extending the 2δ-separation above. Let $\rho_{\mathcal{V}}$ be a metric on the set $\mathcal{V}$, and for $t \geq 0$ define the *neighborhood size* for the set $\mathcal{V}$ by

$$N_t := \max_{\nu \in \mathcal{V}} \mathrm{card}\left\{\nu' \in \mathcal{V} \mid \rho_{\mathcal{V}}(\nu, \nu') \leq t\right\} \tag{19}$$

and the separation function

$$\delta(t) := \frac{1}{2} \min\left\{\rho(\theta(P_\nu), \theta(P_{\nu'})) \mid \nu, \nu' \in \mathcal{V} \text{ and } \rho_{\mathcal{V}}(\nu, \nu') > t\right\}. \tag{20}$$

Then we have the following generalization [14, Corollary 1] of the Fano bound: for any $t \geq 0$,

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \geq \frac{\Phi(\delta(t))}{2} \left\{1 - \frac{I(V; X_1^n) + \log 2}{\log |\mathcal{V}| - \log N_t}\right\}. \tag{21}$$

4.1 A private version of Fano's method

We now turn to developing a version of Fano's lower bound that applies to estimators $\widehat{\theta}$ that act on privatized samples $Z_1^n = (Z_1, \ldots, Z_n)$, where the obfuscation channel Q is non-interactive (Figure 2(b)), meaning that Z_i is conditionally independent of $Z_{\setminus i}$ given X_i. Our upper bound is

variational: it involves optimization over a subset of the space $L^\infty(\mathcal{X}) := \left\{f : \mathcal{X} \to \mathbb{R} \mid \|f\|_\infty < \infty\right\}$ of uniformly bounded functions equipped with the supremum norm $\|f\|_\infty = \sup_x |f(x)|$ and the associated 1-ball of the supremum norm

$$\mathbb{B}_\infty(\mathcal{X}) := \{\gamma \in L^\infty(\mathcal{X}) \mid \|\gamma\|_\infty \leq 1\}. \tag{22}$$

As the set $\mathcal{X}$ is generally clear from context, we typically omit this dependence (and adopt the shorthand $\mathbb{B}_\infty$). As with the classical Fano method, we consider a 2δ-separated family of distributions $\{P_\nu, \nu \in \mathcal{V}\}$, and for each $\nu \in \mathcal{V}$, we define the linear functional $\varphi_\nu : L^\infty(\mathcal{X}) \to \mathbb{R}$ by

$$\varphi_\nu(\gamma) = \int_{\mathcal{X}} \gamma(x)(dP_\nu(x) - d\overline{P}(x)). \tag{23}$$

With this notation, we have the following private version of Fano's method:

Proposition 2 (Private Fano method)**.** *Given any set $\{P_\nu, \nu \in \mathcal{V}\}$, for any $t \geq 0$ the non-interactive α-private minimax risk has lower bound*

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, \alpha) \geq \frac{\Phi(\delta(t))}{2} \left\{1 - \frac{n(e^\alpha - 1)^2}{\log(|\mathcal{V}|/N_t)} \left[\frac{1}{|\mathcal{V}|} \sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} (\varphi_\nu(\gamma))^2\right] - \frac{\log 2}{\log(|\mathcal{V}|/N_t)}\right\}.$$

Underlying Proposition 2 is a variational bound on the mutual information between a sequence $Z_1^n = (Z_1, \ldots, Z_n)$ of private random variables and a random index V drawn uniformly on $\mathcal{V}$, where $Z_1^n \sim M_\nu^n$, conditional on $V = \nu$; that is, Z_1^n is marginally drawn according to the mixture

$$\overline{M} := \frac{1}{|\mathcal{V}|} \sum_{\nu \in \mathcal{V}} M_\nu^n \quad \text{where } M_\nu^n(S) = \int Q(S \mid x_1^n) dP_\nu^n(x_1^n).$$

(Recall equation (15)). When the conditional distribution Q is non-interactive, as considered in this section, then M_ν^n is also a product distribution. By comparison with equation (18), we see that $\overline{M}$ is the private analogue of the mixture distribution $\overline{P}$ that arises in the classical Fano analysis.

Proposition 2 is an immediate consequence of the Fano bound (21) coupled with the following upper bound on the mutual information between Z_1^n and an index V uniformly distributed over $\mathcal{V}$:

$$I(V; Z_1^n) \leq n(e^\alpha - 1)^2 \frac{1}{|\mathcal{V}|} \sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} (\varphi_\nu(\gamma))^2. \tag{24}$$

The inequality (24) is in turn an immediate consequence of Theorem 2 to come; we provide the proof of this inequality in Appendix C.2. We conjecture that it also holds in the fully interactive setting, but given well-known difficulties of characterizing multiple channel capacities with feedback [13, Chapter 15], it may be challenging to verify this conjecture.

4.2 Some applications of the private Fano bound

In this section, we show how Proposition 2 leads to sharp characterizations of the α-private minimax rates for some classical and high-dimensional mean estimation problems. We consider estimation of the d-dimensional mean $\theta(P) := \mathbb{E}_P[X]$ of a random vector $X \in \mathbb{R}^d$. Due to the difficulties associated with differential privacy on non-compact spaces (recall Section 3.2.1), we focus on distributions with compact support. We provide proofs of our mean estimation results in Appendix D.

4.2.1 Classical mean estimation in d-dimensions

We begin by considering estimation of means for sampling distributions supported on ℓ_p balls, where $p \in [1, 2]$. Indeed, for a radius $r < \infty$, consider the family

$$\mathcal{P}_{p,r} := \left\{\text{distributions } P \text{ supported on } \mathbb{B}_p(r) \subset \mathbb{R}^d\right\},$$

where $\mathbb{B}_p(r) = \{x \in \mathbb{R}^d \mid \|x\|_p \leq r\}$ is the ℓ_p-ball of radius r. In the non-private setting, the standard mean estimator $\widehat{\theta} = \frac{1}{n}\sum_{i=1}^n X_i$ has mean-squared error at most r^2/n, since $\|X\|_2 \leq \|X\|_p \leq r$ by assumption. The following result shows that the private minimax MSE is substantially different:

Corollary 5. *For the mean estimation problem, for all $p \in [1, 2]$ and privacy levels $\alpha \in [0, 1]$, the non-interactive private minimax risk satisfies*

$$r^2 \min\left\{1, \frac{1}{\sqrt{n\alpha^2}}, \frac{d}{n\alpha^2}\right\} \lesssim \mathfrak{M}_n(\theta(\mathcal{P}_{p,r}), \|\cdot\|_2^2, \alpha) \lesssim r^2 \min\left\{\frac{d}{n\alpha^2}, 1\right\}.$$

See Appendix D.1 for the proof of this claim; the lower bound makes use of the private form of Fano's method (Proposition 2), while the upper bound is a consequence of the optimal mechanisms we develop in Section 4.2.3 presently. Corollary 5 demonstrates the substantial difference between d-dimensional mean estimation in private and non-private settings: the privacy constraint leads to a multiplicative penalty of d/α^2 in terms of mean-squared error. Thus, the effect of privacy is to reduce the effectve sample size from n to $\alpha^2 n/d$.

4.2.2 Estimation of high-dimensional sparse vectors

Recently, there has been substantial interest in high-dimensional problems in which the dimension d is larger than the sample size n, but a low-dimensional latent structure makes inference possible. (See the paper by Negahban et al. [39] for a general overview.) Here we consider a simple but canonical instance of a high-dimensional problem, that of estimating a sparse mean vector. For an integer parameter $s \geq 1$, consider the class of distributions

$$\mathcal{P}^s_{\infty,r} := \left\{\text{distributions } P \text{ supported on } \mathbb{B}_\infty(r) \subset \mathbb{R}^d \text{ with } \|\mathbb{E}_P[X]\|_0 \leq s\right\}. \tag{25}$$

In the non-private case, estimation of such an s-sparse predictor in the squared ℓ_2-norm is possible at rate $\mathbb{E}[\|\widehat{\theta} - \theta\|_2^2] \lesssim r^2 \frac{s \log(d/s)}{n}$, so that the dimension d can be exponentially larger than the sample size n. With this context, the next result shows that local privacy can have a dramatic impact in the high-dimensional setting. For simplicity, we restrict ourselves to the easiest case of a 1-sparse vector ($s = 1$).

Corollary 6. *For the 1-sparse means problem, for all $\alpha \in [0, 1]$, the non-interactive private minimax risk satisfies*

$$\min\left\{r^2, \frac{r^2 d \log(2d)}{n\alpha^2}\right\} \lesssim \mathfrak{M}_n\left(\theta(\mathcal{P}^1_{\infty,r}), \|\cdot\|_2^2, \alpha\right) \lesssim \min\left\{r^2, \frac{r^2 d \log(2d)}{n\alpha^2}\right\}.$$

See Appendix D.2 for a proof.

From the lower bound in Corollary 6, we see that local differential privacy has an especially dramatic effect for the sparse means problem: due to the presence of the d-term in the denominator, estimation in high-dimensional settings ($d \geq n$) becomes impossible, even for 1-sparse vectors. Contrasting this fact with the scaling $d \asymp e^n$ that 1-sparsity allows in the non-private setting shows that local differential privacy is a very severe constraint in this setting.

4.2.3 Optimal mechanisms: attainability for mean estimation

Our lower bounds for both d-dimensional mean estimation (Corollary 5) and 1-sparse mean estimation (Corollary 6) are based on the private form of Fano's method (Proposition 2). On the other hand, the upper bounds are based on direct analysis of specific privacy mechanisms and estimators. Here we discuss the optimal privacy mechanisms for these two problems in more detail.

Sub-optimality of Laplacian mechanism: For the 1-dimensional mean estimation problem (Corollary 1), we showed that adding Laplacian noise to (truncated versions of) the observations led to an optimal privacy mechanism. The extension of this result to the d-dimensional problems considered in Corollary 5, however, *fails*. More concretely, as a special case of the families in Corollary 5, consider the class $\mathcal{P}_{2,1}$ of distributions supported on the Euclidean ball $\mathbb{B}_2(1) \subset \mathbb{R}^d$ of unit norm. In order to guarantee α-differential privacy, suppose that we output the additively corrupted random vector $Z := x+W$, where the noise vector $W \in \mathbb{R}^d$ has i.i.d components following a Laplace$(\alpha/\sqrt{d})$ distribution. With this choice, it can be verified that for X taking values in $\mathbb{B}_2(1)$, the random vector Z is an α-DLP view of X. However, this privacy mechanism does *not* achieve the minimax risk over α-private mechanisms. In particular, one must suffer the rate

$$\inf_{\widehat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\|\widehat{\theta}(Z_1, \ldots, Z_n) - \mathbb{E}_P[X]\|_2^2 \right] \gtrsim \min \left\{ \frac{d^2}{n\alpha^2}, 1 \right\}, \tag{26}$$

a quadratic (d^2) dimension dependence, as opposed to the linear scaling (d) of the optimal result in Corollary 5. See Appendix D.3 for the proof of claim (26). The poorer dimension dependence of the Laplacian mechanism demonstrates that sampling mechanisms must be chosen carefully.

Optimal mechanisms: Let us now describe some mechanisms that are optimal for the d-dimensional and 1-sparse mean estimation problems. Both of them are procedures that output a random variable Z that is an α-differentially-private view of X, and they are unbiased in the sense that $\mathbb{E}[Z \mid X = x] = x$. They require the Bernoulli random variable

$$T \sim \text{Bernoulli}(\pi_\alpha), \quad \text{where} \quad \pi_\alpha := e^\alpha/(e^\alpha + 1).$$

Privacy mechanism for ℓ_2-ball: Given a vector $x \in \mathbb{R}^d$ with $\|x\|_2 \leq r$, define a random vector

$$\widetilde{X} := \begin{cases} +r\frac{x}{\|x\|_2} & \text{with probability } \frac{1}{2} + \frac{\|x\|_2}{2r} \\ -r\frac{x}{\|x\|_2} & \text{with probability } \frac{1}{2} - \frac{\|x\|_2}{2r}. \end{cases}$$

Then sample $T \sim \text{Bernoulli}(\pi_\alpha)$ and set

$$Z \sim \begin{cases} \text{Uniform}(z \in \mathbb{R}^d \mid \langle z, \widetilde{X} \rangle > 0, \|z\|_2 = B) & \text{if } T = 1 \\ \text{Uniform}(z \in \mathbb{R}^d \mid \langle z, \widetilde{X} \rangle \leq 0, \|z\|_2 = B) & \text{if } T = 0, \end{cases} \tag{27}$$

where B is chosen to equal

$$B = r\frac{e^\alpha + 1}{e^\alpha - 1}\frac{\sqrt{\pi}}{2}\frac{d\,\Gamma(\frac{d-1}{2} + 1)}{\Gamma(\frac{d}{2} + 1)}.$$

Privacy mechanism for ℓ_∞-ball: Given a vector $x \in \mathbb{R}^d$ with $\|x\|_\infty \le r$, construct a random vector $\widetilde{X} \in \mathbb{R}^d$ with independent coordinates of the form

$$\widetilde{X}_j = \begin{cases} +r & \text{with probablity } \frac{1}{2} + \frac{x_j}{2r} \\ -r & \text{with probablity } \frac{1}{2} - \frac{x_j}{2r}. \end{cases}$$

Then sample $T \sim \text{Bernoulli}(\pi_\alpha)$ and set

$$Z \sim \begin{cases} \text{Uniform}(z \in \{-B, B\}^d \mid \langle z, \widetilde{X} \rangle \ge 0) & \text{if } T = 1 \\ \text{Uniform}(z \in \{-B, B\}^d \mid \langle z, \widetilde{X} \rangle \le 0) & \text{if } T = 0, \end{cases} \tag{28}$$

where the value B is chosen to equal

$$B = r\frac{e^\alpha + 1}{e^\alpha - 1} C_d, \quad \text{where} \quad C_d^{-1} = \begin{cases} \frac{1}{2^{d-1}} \binom{d-1}{(d-1)/2} & \text{if } d \text{ is odd} \\ \frac{1}{2^{d-1} + \frac{1}{2}\binom{d}{d/2}} \binom{d-1}{d/2} & \text{if } d \text{ is even.} \end{cases}$$

See Figure 4 for visualizations of the geometry that underlies these strategies. By construction, each scheme guarantees that Z is an α-private view of X. Each strategy is efficiently implementable when combined with rejection sampling: the ℓ_2-mechanism (27) by normalizing a sample from the $\mathsf{N}(0, I_{d\times d})$ distribution, and the ℓ_∞-strategy (28) by sampling the hypercube $\{-1, 1\}^d$. Additionally, by Stirling's approximation, we have that in each case $B \lesssim \frac{r\sqrt{d}}{\alpha}$ for $\alpha \le 1$. Moreover, they are unbiased (see Appendix I.2 for the unbiasedness of strategy (27) and Appendix I.3 for strategy (28)).

We now complete the picture of minimax optimal estimation schemes for Corollaries 5 and 6. In the case of Corollary 5, the estimator $\widehat{\theta} := \frac{1}{n}\sum_{i=1}^n Z_i$, where Z_i is constructed by procedure (27) from the non-private vector X_i, attains minimax optimal convergence rate. In the case of Corollary 6, a slightly more complex estimator gives the result: in particular, we set

$$\widehat{\theta}_n := \operatorname*{argmin}_{\theta \in \mathbb{R}^d} \left\{ \frac{1}{2n} \sum_{i=1}^n \|Z_i - \theta\|_2^2 + \lambda_n \|\theta\|_1 \right\}, \quad \text{where} \quad \lambda_n = 2\frac{\sqrt{d \log d}}{\sqrt{n\alpha^2}},$$

and the Z_i are drawn according to strategy (28). See Appendix D.2 for a rigorous argument.

4.3 Variational bounds on private mutual information

The private Fano bound in Proposition 2 reposes on a variational bound on the private mutual information, which we describe here. Recall the space $L^\infty(\mathcal{X}) := \{f : \mathcal{X} \to \mathbb{R} \mid \|f\|_\infty < \infty\}$ of uniformly bounded functions, equipped with the usual sup-norm and unit norm ball (22), $\mathbb{B}_\infty$, as well as the linear functionals φ_ν from Eq. (23). We then have the following result.

Theorem 2. *Let $\{P_\nu\}_{\nu\in\mathcal{V}}$ be an arbitrary collection of probability measures on $\mathcal{X}$, and let $\{M_\nu\}_{\nu\in\mathcal{V}}$ be the set of marginal distributions induced by an α-differentially private distribution Q. Then*

$$\frac{1}{|\mathcal{V}|} \sum_{\nu\in\mathcal{V}} \left[D_{\text{kl}}\left(M_\nu \| \overline{M}\right) + D_{\text{kl}}\left(\overline{M} \| M_\nu\right) \right] \le \frac{(e^\alpha - 1)^2}{|\mathcal{V}|} \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{\nu\in\mathcal{V}} \left(\varphi_\nu(\gamma)\right)^2.$$

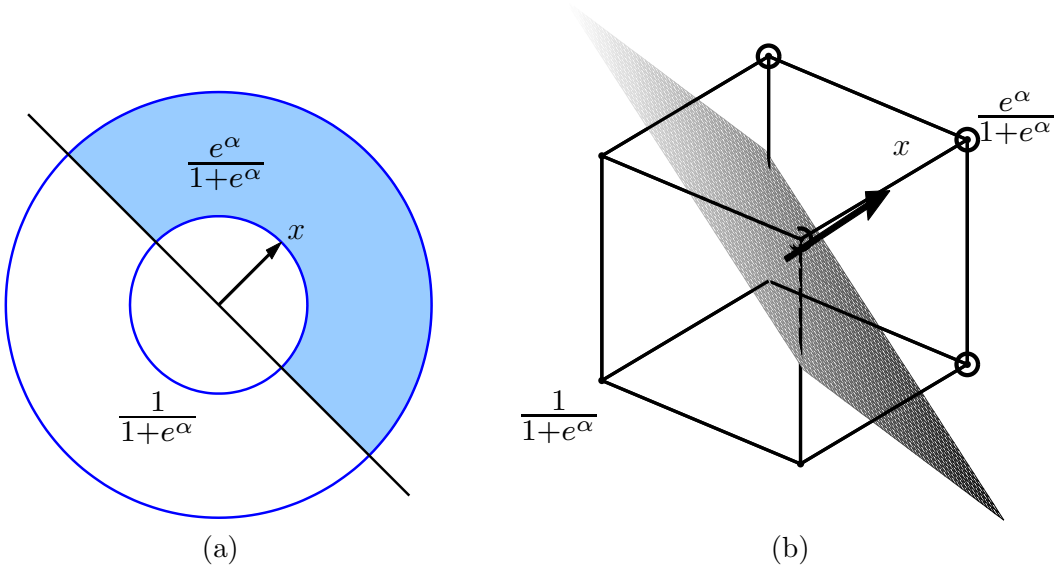

Figure 4. Private sampling strategies for different types of ℓ_p-balls. (a) A private sampling strategy (27) for data lying within an ℓ_2-ball. Outer boundary of highlighted region sampled uniformly with probability $e^\alpha/(e^\alpha+1)$. (b) A private sampling strategy (28) for data lying within an ℓ_∞-ball. Circled point set sampled uniformly with probability $e^\alpha/(e^\alpha+1)$.

It is important to note that, at least up to constant factors, Theorem 2 is never weaker than the results provided by Theorem 1. By definition of the linear functional φ_ν, we have

$$\sup_{\gamma\in\mathbb{B}_\infty(\mathcal{X})}\sum_{\nu\in\mathcal{V}}\left(\varphi_\nu(\gamma)\right)^2 \overset{(i)}{\leq} \sum_{\nu\in\mathcal{V}}\sup_{\gamma\in\mathbb{B}_\infty(\mathcal{X})}\left(\varphi_\nu(\gamma)\right)^2 = 4\sum_{\nu\in\mathcal{V}}\left\|P_\nu-\overline{P}\right\|_{\mathrm{TV}}^2,$$

where inequality (i) follows by interchanging the summation and supremum. Overall, we have

$$I(Z;V)=\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}D_{\mathrm{kl}}\left(M_\nu\|\overline{M}\right)\leq 4(e^\alpha-1)^2\frac{1}{|\mathcal{V}|^2}\sum_{\nu,\nu'\in\mathcal{V}}\|P_\nu-P_{\nu'}\|_{\mathrm{TV}}^2.$$

The strength of Theorem 2 arises from the fact that inequality (i)—the interchange of the order of supremum and summation—may be quite loose.

We may extend Theorem 2 to sequences of random variables, that is, to collections $\{P_\nu^n\}_{\nu\in\mathcal{V}}$ of product probability measures in the non-interactive case. Indeed, we have by the standard chain rule for mutual information [26, Chapter 5] that

$$I(Z_1,\ldots,Z_n;V)=\sum_{i=1}^n I(Z_i;V\mid Z_{1:i-1})\leq\sum_{i=1}^n I(Z_i;V),$$

where the inequality is a consequence of the conditional independence of the variables Z_i given V, which holds when the channel Q is non-interactive. Applying Theorem 2 to the individual terms $I(Z_i;V)$ then yields inequality (24); see Appendix C.2 for a fully rigorous derivation.

5 Bounds on multiple pairwise divergences: Assouad's method

Thus far, we have seen how Le Cam's method and Fano's method, in the form of Propositions 1 and 2, can be used to derive sharp minimax rates. However, their application appears to be limited to problems whose minimax rates can be controlled via reductions to binary hypothesis tests (Le Cam's method) or for non-interactive privacy mechanisms (Fano's method). Another classical approach to deriving minimax lower bounds is Assouad's method [4, 52]. In this section, we show that a privatized form of Assouad's method can be be used to obtain sharp minimax rates in interactive settings. We illustrate by deriving bounds for several problems, including multinomial probability estimation and nonparametric density estimation.

Assouad's method transforms an estimation problem into multiple binary hypothesis testing problems, using the structure of the problem in an essential way. For some $d \in \mathbb{N}$, let $\mathcal{V} = \{-1, 1\}^d$, and let us consider a family of distributions $\{P_\nu\}_{\nu \in \mathcal{V}}$ indexed by the hypercube. We say that the the family $\{P_\nu\}_{\nu \in \mathcal{V}}$ induces a 2δ-*Hamming separation* for the loss $\Phi \circ \rho$ if there exists a vertex mapping (a function $\mathsf{v} : \theta(\mathcal{P}) \to \{-1, 1\}^d$) satisfying

$$\Phi(\rho(\theta, \theta(P_\nu))) \geq 2\delta \sum_{j=1}^{d} \mathbf{1}\{[\mathsf{v}(\theta)]_j \neq \nu_j\}. \tag{29}$$

As in the standard reduction from estimation to testing, we consider the following random process: nature chooses a vector $V \in \{-1, 1\}^d$ uniformly at random, after which the sample $X_1, \ldots, X_n$ is drawn from the distribution P_ν conditional on $V = \nu$. Letting $\mathbb{P}_{\pm j}$ denote the joint distribution over the random index V and X conditional on the jth coordinate $V_j = \pm 1$, we obtain the following sharper variant of Assouad's lemma [4].

Lemma 1 (Sharper Assouad method)**.** *Under the conditions of the previous paragraph, we have*

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \geq \delta \sum_{j=1}^{d} \inf_{\psi} \left[\mathbb{P}_{+j}(\psi(X_{1:n}) \neq +1) + \mathbb{P}_{-j}(\psi(X_{1:n}) \neq -1)\right].$$

We provide a proof of Lemma 1 in Section I.1 (see also the paper [3]). We can also give a variant of Lemma 1 after some minor rewriting. For each $j \in [d]$ define the mixture distributions

$$P^n_{+j} = \frac{1}{2^{d-1}} \sum_{\nu : \nu_j = 1} P^n_\nu, \quad P^n_{-j} = \frac{1}{2^{d-1}} \sum_{\nu : \nu_j = -1} P^n_\nu, \quad P_{\pm j} = \frac{1}{2^{d-1}} \sum_{\nu : \nu_j = \pm 1} P_\nu \tag{30}$$

where P^n_ν is the (product) distribution of $X_1, \ldots, X_n$. Then, by Le Cam's lemma, the following minimax lower bound is equivalent to the Assouad bound of Lemma 1:

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho) \geq \delta \sum_{j=1}^{d} \left[1 - \left\|P^n_{+j} - P^n_{-j}\right\|_{\mathrm{TV}}\right]. \tag{31}$$

5.1 A private version of Assoud's method

As in the preceding sections, we extend Lemma 1 to the locally differentially private setting. In this case, we are able to provide a minimax lower bound that applies to any locally differentially

private channel Q, including in interactive settings (Fig. 2(a)). In this case, we again let $\mathbb{B}_\infty(\mathcal{X})$ denote collection of functions $f : \mathcal{X} \to \mathbb{R}$ with supremum norm bounded by 1 (definition (22)). Then we have the following private version of Assouad's method.

Proposition 3 (Private Assouad bound)**.** *Let the conditions of Lemma 1 hold, that is, let the family $\{P_\nu\}_{\nu \in \mathcal{V}}$ induce a 2δ-Hamming separation for the loss $\Phi \circ \rho$. Then*

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, \alpha) \geq d\delta \left[1 - \left(\frac{n(e^\alpha - 1)^2}{2d} \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{j=1}^{d} \left(\int_{\mathcal{X}} \gamma(x)(dP_{+j}(x) - dP_{-j}(x))\right)^2\right)^{\frac{1}{2}}\right].$$

As is the case for our private analogue of Fano's method (Proposition 2), underlying Proposition 3 is a variational bound that generalizes the usual total variation distance to a variational quantity applied jointly to multiple mixtures of distributions. Proposition 3 is an immediate consequence of Theorem 3 to come.

5.2 Some applications of the private Assouad bound

Proposition 3 gives sharp characterizations of α-private minimax rates for a number of classical statistical problems. While it requires that there be a natural coordinate-wise structure to the problem at hand because of the Hamming separation condition (29), such conditions are common in a number of estimation problems. Additionally, Proposition 3 applies to interactive channels Q. As examples, we consider multinomial probability and non-parametric density estimation.

5.2.1 Multinomial estimation under local privacy

For our first application of Proposition 3, we return to the original motivation for local privacy [49]: avoiding answer bias in surveys. Consider the probability simplex

$$\Delta_d := \left\{\theta \in \mathbb{R}^d \mid \theta \geq 0 \text{ and } \sum_{j=1}^{d} \theta_j = 1\right\}.$$

Any vector $\theta \in \Delta_d$ specifies a multinomial random variable taking d states, in particular with probabilities $P_\theta(X = j) = \theta_j$ for $j \in \{1, \ldots, d\}$. Given a sample from this distribution, our goal is to estimate the probability vector θ. Warner [49] studied the Bernoulli variant of this problem (corresponding to $d = 2$), proposing a mechanism known as *randomized response*: For a given survey question, respondents answer truthfully with probability $p > 1/2$ and lie with probability $1 - p$. Here we show that an extension of this mechanism is optimal for α-locally differentially private multinomial estimation via the following corollary to Proposition 3, which provides bounds on the minimax error in both squared ℓ_2-norm and ℓ_1-norm for (sequentially) interactive channels.

Corollary 7. *For the multinomial estimation problem, for any α-locally differentially private channel (3), there exist universal constants $0 < c_\ell \leq c_u < 5$ such that for all $\alpha \in [0, 1]$,*

$$c_\ell \min\left\{1, \frac{1}{\sqrt{n\alpha^2}}, \frac{d}{n\alpha^2}\right\} \leq \mathfrak{M}_n\left(\Delta_d, \|\cdot\|_2^2, \alpha\right) \leq c_u \min\left\{1, \frac{d}{n\alpha^2}\right\}, \quad \text{and} \tag{32a}$$

$$c_\ell \min\left\{1, \frac{d}{\sqrt{n\alpha^2}}\right\} \leq \mathfrak{M}_n\left(\Delta_d, \|\cdot\|_1, \alpha\right) \leq c_u \min\left\{1, \frac{d}{\sqrt{n\alpha^2}}\right\}. \tag{32b}$$

See Appendix F for the proofs of the lower bounds. We provide simple estimation strategies achieving the upper bounds presently.

Before presenting the (minimax) optimal estimation scheme, we compare the private rates to classical ones without privacy. The maximum likelihood estimate sets $\widehat{\theta}_j$ as the proportion of the sample taking value j; it has mean-squared error

$$\mathbb{E}\left[\|\widehat{\theta} - \theta\|_2^2\right] = \sum_{j=1}^{d} \mathbb{E}\left[(\widehat{\theta}_j - \theta_j)^2\right] = \frac{1}{n}\sum_{j=1}^{d} \theta_j(1-\theta_j) \le \frac{1}{n}\left(1 - \frac{1}{d}\right) < \frac{1}{n}.$$

An analogous calculation for the ℓ_1-norm yields $\mathbb{E}[\|\widehat{\theta} - \theta\|_1] \le \sqrt{d}\mathbb{E}[\|\widehat{\theta} - \theta\|_2] < \sqrt{d/n}$. Consequently, for estimation in ℓ_1 or ℓ_2-norm, the effect of providing α-differential privacy causes the effective sample size to decrease as $n \mapsto n\alpha^2/d$.

A consequence of the lower bound (32a) is the following: a minor variant of Warner's randomized response strategy is an optimal mechanism, as we now show. Since there is a bijection between multinomial observations $x \in \{1, \ldots, d\}$ and the d standard basis vectors $e_1, \ldots, e_d \in \mathbb{R}^d$, we abuse notation and represent observations as either as convenient. Our (extended) randomized response procedure is as follows: we construct the private vector $Z \in \{0,1\}^d$ from a multinomial observation $x \in \{e_1, \ldots, e_d\}$ by sampling d coordinates independently via the procedure

$$[Z]_j = \begin{cases} x_j & \text{with probability } \frac{\exp(\alpha/2)}{1+\exp(\alpha/2)} \\ 1 - x_j & \text{with probability } \frac{1}{1+\exp(\alpha/2)}. \end{cases} \tag{33}$$

The distribution (33) is α-differentially private: for $x, x' \in \Delta_d$ and any $z \in \{0,1\}^d$, we have

$$\frac{Q(Z = z \mid x)}{Q(Z = z \mid x')} = \exp\left(\frac{\alpha}{2}\left(\|z - x\|_1 - \|z - x'\|_1\right)\right) \in \left[\exp(-\alpha), \exp(\alpha)\right],$$

where the triangle inequality guarantees $|\, \|z - x\|_1 - \|z - x'\|_1 \,| \le 2$. We now compute the expected value and variance of the random variables Z. Using the definition (33), we have

$$\mathbb{E}[Z \mid x] = \frac{e^{\alpha/2}}{1 + e^{\alpha/2}} x + \frac{1}{1 + e^{\alpha/2}}(\mathbb{1} - x) = \frac{e^{\alpha/2} - 1}{e^{\alpha/2} + 1} x + \frac{1}{1 + e^{\alpha/2}}\mathbb{1}.$$

As the random variables Z are Bernoulli, we have the variance bound $\mathbb{E}[\|Z\|_2^2] \le d/4$. Letting Π_{Δ_d} denote the projection operator onto the simplex, we arrive at the natural estimator

$$\widehat{\theta}_{\text{part}} := \frac{1}{n}\sum_{i=1}^{n}\left(Z_i - \mathbb{1}/(1 + e^{\alpha/2})\right)\frac{e^{\alpha/2} + 1}{e^{\alpha/2} - 1} \quad \text{and} \quad \widehat{\theta} := \Pi_{\Delta_d}\left(\widehat{\theta}_{\text{part}}\right). \tag{34}$$

The projection of $\widehat{\theta}_{\text{part}}$ onto the probability simplex can be done in time linear in the dimension d of the problem [9], so the estimator (34) is efficiently computable. Since projections onto convex sets are non-expansive, any pair of vectors in the simplex are at most ℓ_2-distance $\sqrt{2}$ apart, and $\mathbb{E}_\theta[\widehat{\theta}_{\text{part}}] = \theta$ by construction, we have for $\alpha \le 1$ that

$$\mathbb{E}\left[\|\widehat{\theta} - \theta\|_2^2\right] \le \min\left\{2, \mathbb{E}\left[\|\widehat{\theta}_{\text{part}} - \theta\|_2^2\right]\right\} \le \min\left\{2, \frac{d}{4n}\left(\frac{e^{\alpha/2} + 1}{e^{\alpha/2} - 1}\right)^2\right\} \le \min\left\{2, 5\frac{d}{n\alpha^2}\right\}.$$

Similar results hold for the ℓ_1-norm with the same estimator: by Cauchy-Schwarz and Jensen's inequality, $\mathbb{E}[\|\widehat{\theta} - \theta\|_1] \leq \min\{1, \sqrt{d}\mathbb{E}[\|\widehat{\theta} - \theta\|_2]\} \leq \min\{1, \sqrt{5}d/\sqrt{n\alpha^2}\}$.

There are also other relatively simple estimation strategies that achieve convergence rate $d/n\alpha^2$; the Laplace perturbation approach [19] is another. Nonetheless, its ease of use, coupled with our optimality results, provide support for randomized response as a desirable mechanism for private probability estimation. See also our experiments in Section 6.3 on logistic estimation, which show that randomized response generally yields better performance than Laplace noise addition.

5.2.2 Density estimation under local privacy

In this section, we show that the effects of local differential privacy are more severe for nonparametric density estimation: instead of just a multiplicative loss in the effective sample size as in previous sections, imposing local differential privacy leads to a different convergence rate. This result holds even though we solve a problem in which both the function being estimated and the observations themselves belong to compact spaces.

Definition 3 (Elliptical Sobolev space)**.** For a given orthonormal basis $\{\varphi_j\}$ of $L^2([0,1])$, smoothness parameter $\beta > 1/2$ and radius r, the Sobolev class of order β is given by

$$\mathcal{F}_\beta[r] := \left\{ f \in L^2([0,1]) \mid f = \sum_{j=1}^{\infty} \theta_j \varphi_j \text{ such that } \sum_{j=1}^{\infty} j^{2\beta}\theta_j^2 \leq r^2 \right\}.$$

If we choose the trignometric basis as our orthonormal basis, membership in the class $\mathcal{F}_\beta[r]$ corresponds to smoothness constraints on the derivatives of f. More precisely, for $j \in \mathbb{N}$, consider the orthonormal basis for $L^2([0,1])$ of trigonometric functions:

$$\varphi_0(t) = 1, \quad \varphi_{2j}(t) = \sqrt{2}\cos(2\pi jt), \quad \varphi_{2j+1}(t) = \sqrt{2}\sin(2\pi jt). \tag{35}$$

Let f be a β-times almost everywhere differentiable function for which $|f^{(\beta)}(x)| \leq r$ for almost every $x \in [0,1]$ satisfying $f^{(k)}(0) = f^{(k)}(1)$ for $k \leq \beta - 1$. Then, uniformly over all such f, there is a universal constant $c \leq 2$ such that that $f \in \mathcal{F}_\beta[cr]$ (see, for instance, [47, Lemma A.3]).

Suppose our goal is to estimate a density function $f \in \mathcal{F}_\beta[C]$ and that quality is measured in terms of the squared error (squared $L^2[0,1]$-norm)

$$\|\widehat{f} - f\|_2^2 := \int_0^1 (\widehat{f}(x) - f(x))^2 dx.$$

The well-known [52, 51, 47] (non-private) minimax mean squared error scales as

$$\mathfrak{M}_n\left(\mathcal{F}_\beta, \|\cdot\|_2^2, \infty\right) \asymp n^{-\frac{2\beta}{2\beta+1}}. \tag{36}$$

The goal of this section is to understand how this minimax rate changes when we add an α-privacy constraint to the problem. Our main result is to demonstrate that the classical rate (36) is no longer attainable when we require α-local differential privacy.

Lower bounds on density estimation We begin by giving our main lower bound on the minimax rate of estimation of densities when observations from the density are differentially private. We provide the proof of the following result in Appendix G.

Corollary 8. *For the class of densities $\mathcal{F}_\beta$ defined using the trigonometric basis (35), there exist constants $0 < c_\beta \leq c'_\beta < \infty$, dependent only on β, such that the α-private minimax risk (for $\alpha \in [0, 1]$) is sandwiched as*

$$c_\beta \left(n\alpha^2\right)^{-\frac{2\beta}{2\beta+2}} \leq \mathfrak{M}_n \left(\mathcal{F}_\beta[1], \|\cdot\|_2^2, \alpha\right) \leq c'_\beta \left(n\alpha^2\right)^{-\frac{2\beta}{2\beta+2}}. \tag{37}$$

The most important feature of the lower bound (37) is that it involves a *different polynomial exponent* than the classical minimax rate (36). Whereas the exponent in classical case (36) is $2\beta/(2\beta+1)$, it reduces to $2\beta/(2\beta+2)$ in the locally private setting. For example, when we estimate Lipschitz densities ($\beta = 1$), the rate degrades from $n^{-2/3}$ to $n^{-1/2}$.

Interestingly, no estimator based on Laplace (or exponential) perturbation of the observations X_i themselves can attain the rate of convergence (37). This fact follows from results of Carroll and Hall [10] on nonparametric deconvolution. They show that if observations X_i are perturbed by additive noise W, where the characteristic function ϕ_W of the additive noise has tails behaving as $|\phi_W(t)| = \mathcal{O}(|t|^{-a})$ for some $a > 0$, then no estimator can deconvolve $X + W$ and attain a rate of convergence better than $n^{-2\beta/(2\beta+2a+1)}$. Since the characteristic function of the Laplace distribution has tails decaying as t^{-2}, no estimator based on the Laplace mechanism (applied directly to the observations) can attain rate of convergence better than $n^{-2\beta/(2\beta+5)}$. In order to attain the lower bound (37), we must thus study alternative privacy mechanisms.

Achievability by orthogonal projection estimators For $\beta = 1$, histogram estimators with counts perturbed by Laplacian noise achieve the optimal rate of convergence (37); this is a consequence of the results of Wasserman and Zhou [50, Section 4.2] applied to locally private mechanisms. For higher degrees of smoothness ($\beta > 1$), standard histogram estimators no longer achieve optimal rates in the classical setting [44]. Accordingly, we now turn to developing estimators based on orthogonal series expansion, and show that even in the setting of local privacy, they can achieve the lower bound (37) for all orders of smoothness $\beta \geq 1$.

Recall the elliptical Sobolev space (Definition 3), in which a function f is represented in terms of its basis expansion $f = \sum_{j=1}^\infty \theta_j \varphi_j$. This representation underlies the following classical orthonormal series estimator. Given a sample $X_{1:n}$ drawn i.i.d. according to a density $f \in L^2([0, 1])$, compute the empirical basis coefficients

$$\widehat{\theta}_j = \frac{1}{n} \sum_{i=1}^n \varphi_j(X_i) \quad \text{for } j \in \{1, \ldots, k\}, \tag{38}$$

where the value $k \in \mathbb{N}$ is chosen either a priori based on known properties of the estimation problem or adaptively, for example, using cross-validation [21, 47]. Using these empirical coefficients, the density estimate is $\widehat{f} = \sum_{j=1}^k \widehat{\theta}_j \varphi_j$.

In our local privacy setting, we consider a mechanism that employs a random vector $Z_i = (Z_{i,1}, \ldots, Z_{i,k})$ satisfying the unbiasedness condition $\mathbb{E}[Z_{i,j} \mid X_i] = \varphi_j(X_i)$ for each $j \in [k]$. We assume the basis functions are B_0-uniformly bounded, that is, $\sup_j \sup_x |\varphi_j(x)| \leq B_0 < \infty$. This boundedness condition holds for many standard bases, including the trigonometric basis (35) that

underlies the classical Sobolev classes and the Walsh basis. We generate the random variables from the vector $v \in \mathbb{R}^k$ defined by $v_j = \varphi_j(X)$ in the hypercube-based sampling scheme (28), where we assume $B > B_0$. With this sampling strategy, iteration of expectation yields

$$\mathbb{E}[[Z]_j \mid X = x] = c_k \frac{B}{B_0\sqrt{k}} \left(\frac{e^\alpha}{e^\alpha + 1} - \frac{1}{e^\alpha + 1} \right) \varphi_j(x), \tag{39}$$

where $c_k > 0$ is a constant (which is bounded independently of k). Consequently, it suffices to take $B = \mathcal{O}(B_0\sqrt{k}/\alpha)$ to guarantee the unbiasedness condition $\mathbb{E}[[Z_i]_j \mid X_i] = \varphi_j(X_i)$.

Overall, the privacy mechanism and estimator perform the following steps:

(i) given a data point X_i, set the vector $v = [\varphi_j(X_i)]_{j=1}^k$;

(ii) sample Z_i according to the strategy (28), starting from the vector v and using the bound $B = B_0\sqrt{k}(e^\alpha + 1)/c_k(e^\alpha - 1)$;

(iii) compute the density estimate

$$\widehat{f} := \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^k Z_{i,j} \varphi_j. \tag{40}$$

Whenever the underlying function f belongs to the Sobolev space $\mathcal{F}_\beta[r]$ and the orthonormal basis functions φ_j are uniformly bounded by B_0, then the estimator (40) with the choice $k = (n\alpha^2)^{1/(2\beta+2)}$ has mean squared error upper bounded as

$$\mathbb{E}_f \left[\|f - \widehat{f}\|_2^2 \right] \leq c_{B_0,\beta} \left(n\alpha^2 \right)^{-\frac{2\beta}{2\beta+2}}.$$

This shows that the minimax bound (37) is indeed sharp, and there exist easy-to-compute estimators achieving the guarantee. See Section G.2 for a proof of this inequality.

Before concluding our exposition, we make a few remarks on other potential density estimators. Our orthogonal series estimator (40) and sampling scheme (39), while similar in spirit to that proposed by Wasserman and Zhou [50, Sec. 6], is different in that it is locally private and requires a different noise strategy to obtain both α-local privacy and the optimal convergence rate. Lastly, similarly to our remarks on the insufficiency of standard Laplace noise addition for mean estimation, it is worth noting that density estimators that are based on orthogonal series and Laplace perturbation are sub-optimal: they can achieve (at best) rates of $(n\alpha^2)^{-\frac{2\beta}{2\beta+3}}$. This rate is polynomially worse than the sharp result provided by Corollary 8. Again, we see that appropriately chosen noise mechanisms are crucial for obtaining optimal results.

5.3 Variational bounds on paired divergences

In this section, we provide a theorem that underpins the variational inequality in the minimax lower bound of Proposition 3. We provide a slightly more general bound than that required for Proposition 3, showing how it implies the earlier result. Recall that for some $d \in \mathbb{N}$, we consider collections of distributions indexed using the Boolean hypercube $\mathcal{V} = \{-1, 1\}^d$. For each $i \in [n]$ and $\nu \in \mathcal{V}$, let the distribution $P_{\nu,i}$ be supported on the fixed set $\mathcal{X}$, and define the product distribution

$P_\nu^n = \prod_{i=1}^n P_{\nu,i}$. Then in addition to the definition (30) of the paired mixtures $P_{\pm j}^n$, for each i we let

$$P_{+j,i} = \frac{1}{2^{d-1}} \sum_{\nu:\nu_j=1} P_{\nu,i} \quad \text{and} \quad P_{-j,i} = \frac{1}{2^{d-1}} \sum_{\nu:\nu_j=-1} P_{\nu,i},$$

and in analogy to the marginal channel (15), we define the marginal mixtures

$$M_{+j}^n(S) := \frac{1}{2^{d-1}} \sum_{\nu:\nu_j=1} M_\nu^n(S) = \int Q^n(S \mid x_{1:n}) dP_{+j}^n(x_1, \ldots, x_n) \quad \text{for } j = 1, \ldots, d,$$

with the distributions M_{-j}^n defined similarly. For a given pair of distributions (M, M'), we let $D_{\mathrm{kl}}^{\mathrm{sy}}(M \| M') = D_{\mathrm{kl}}(M \| M') + D_{\mathrm{kl}}(M' \| M)$ denote the symmetrized KL-divergence. Recalling the supremum norm ball $\mathbb{B}_\infty(\mathcal{X}) = \{f : \mathcal{X} \to \mathbb{R} \mid \|f\|_\infty \le 1\}$ of Eq. (22), we have the following theorem.

Theorem 3. *Under the conditions of the previous paragraph, for any α-locally differentially private (3) channel Q, we have*

$$\sum_{j=1}^d D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j}^n \| M_{-j}^n\right) \le 2(e^\alpha - 1)^2 \sum_{i=1}^n \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{j=1}^d \left(\int_{\mathcal{X}} \gamma(x) \left(dP_{+j,i}(x) - dP_{-j,i}(x)\right) \right)^2.$$

Theorem 3 generalizes Theorem 1, which corresponds to the special case $d = 1$, though it also has parallels with Theorem 2, as taking the supremum outside the summation is essential to obtain sharp results. We provide the proof of Theorem 3 in Section E.

Theorem 3 allows us to prove Proposition 3 from the non-private variant (31) of Assouad's method. Indeed, inequality (31) immediately implies for any channel Q that

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, Q) \ge \delta \sum_{j=1}^d \left[1 - \left\|M_{+j}^n - M_{-j}^n\right\|_{\mathrm{TV}}\right].$$

Using a combination of Pinsker's inequality and Cauchy-Schwarz, we obtain

$$\sum_{j=1}^d \left\|M_{+j}^n - M_{-j}^n\right\|_{\mathrm{TV}} \le \frac{1}{2}\sqrt{d}\left(\sum_{j=1}^d D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\mathrm{kl}}\left(M_{-j}^n \| M_{+j}^n\right) \right)^{\frac{1}{2}}.$$

Thus, whenever P_ν induces a 2δ-Hamming separation for $\Phi \circ \rho$ we have

$$\mathfrak{M}_n(\theta(\mathcal{P}), \Phi \circ \rho, Q) \ge d\delta \left[1 - \left(\frac{1}{4d} \sum_{j=1}^d D_{\mathrm{kl}}^{\mathrm{sy}}\left(M_{+j}^n \| M_{-j}^n\right) \right)^{\frac{1}{2}}\right]. \tag{41}$$

The combination of inequality (41) with Theorem 3 yields Proposition 3.

6 Experiments

In this section, we study three different datasets—each consisting of data that are sensitive but are nonetheless available publicly—in an effort to demonstrate the importance of minimax theory for practical private estimation. While the public availability of these datasets in some sense obviates the need for private analysis, they provide natural proxies by which to evaluate the performance of privacy-preserving mechanisms and estimation schemes. (The public availability also allows us to make all of our experimental data freely available.)

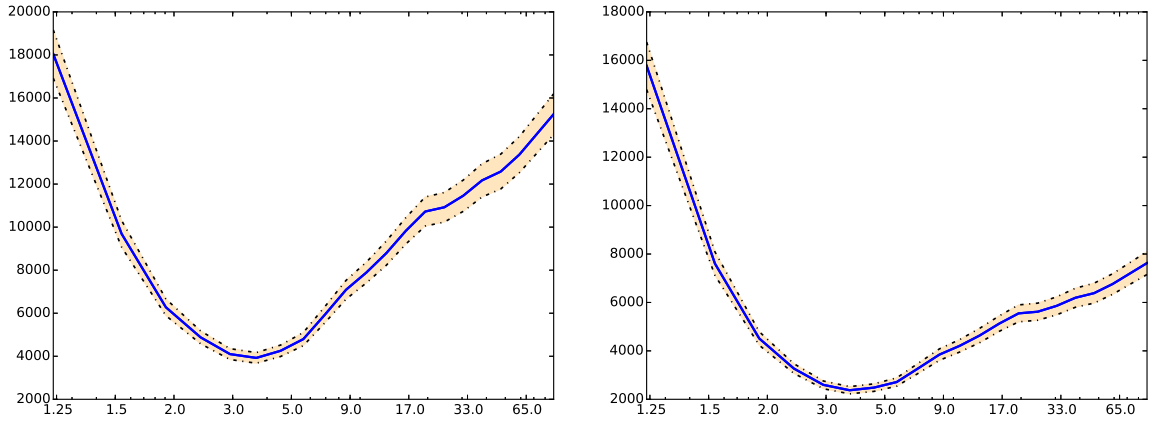

(a) (b)

Figure 5. Mean salary estimation errors. Left (a): privacy $\alpha = 1/2$, right (b): privacy $\alpha = 1$. The horizontal axis corresponds to known assumed moment power k, that is, $\mathbb{E}[|X|^k]^{1/k} = r_k$, and vertical axis is mean absolute error. Confidence bands are simulated 90% confidence intervals.

6.1 Salary estimation: experiments with 1-dimensional data

Our first set of experiments investigates the performance of minimax optimal estimators of one-dimensional mean and median statistics, as described in Sections 3.2.1 and 3.2.2. We use data on salaries in the University of California system [41]. We perform our experiments with the 2010 salaries, which consists of a population of $N = 252{,}540$ employees with mean salary \$39,531 and median salary \$24,968. The data has reasonably long tails; 14 members of the dataset have salaries above \$1,000,000 and two individuals have salaries between two and three-million dollars.

6.1.1 Mean salary estimation

We first explore the effect of bounds on moments in the mean estimation problem as discussed in Section 3.2.1. Recall that for the family of distributions P satisfying $\mathbb{E}_P[|X|^k]^{1/k} \le r_k$, a minimax optimal estimator is the truncated mean estimator (10) (cf. Sec. B.1) with truncation level $T_k = r_k(n\alpha^2)^{\frac{1}{2k}}$; then $\widehat{\theta} = \frac{1}{n}\sum_{i=1}^n Z_i$ has convergence rate $\mathbb{E}[(\widehat{\theta} - \mathbb{E}[X])^2] \lesssim r_k \left(n\alpha^2\right)^{-\frac{k-1}{k}}$. For heavy-tailed data, it may be the case that $\mathbb{E}[|X|^{k_1}]^{1/k_1} \ll \mathbb{E}[|X|^{k_2}]^{1/k_2}$ for some $k_1 < k_2$; thus, while the rate $n^{-\frac{k-1}{k}}$ may be slower for $k = k_1$ than $k = k_2$, the leading moment-based term r_{k_1} may yield better finite sample performance. We investigate this possibility in several experiments.

We perform each experiment by randomly subsampling one-half of the dataset to generate samples $\{X_1, \ldots, X_{N/2}\}$, whose means we estimate privately using the estimator (10). Within an experiment, we assume we have knowledge of the population moment $r_k = \mathbb{E}[|X|^k]^{1/k}$ for large enough k—an unrealistic assumption facilitating comparison. Fixing $m = 20$, we use the $m - 1$ logarithmically-spaced powers

$$k \in \left\{10^{\frac{j}{m-1}\log_{10}(3m)} : j = 1, 2, \ldots, m-1\right\} \approx \{1.24, 1.54, 1.91, \ldots, 60\} \quad \text{and} \quad k = \infty.$$

We repeat the experiment 400 times for each moment value k and privacy parameters $\alpha \in \{.1, .5, 1, 2\}$.

Table 1. Optimal mean absolute erros for each privacy level α, as determined by mean absolute error over assumed known moment k.

Privacy α	$1/10$	$1/2$	1	2	∞
Mean absolute error	11849	3923	2373	1439	82
Moment k	2.9	3.6	3.6	3.6	N/A
Standard error	428	150	91	55	6.0

In Figure 5, we present the behavior of the truncated mean estimator for estimating the mean salary paid in fiscal year 2010 in the UC system. We plot the mean absolute error of the private estimator for the population mean against the moment k for the experiments with $\alpha = \frac{1}{2}$ and $\alpha = 1$. We see that for this (somewhat) heavy-tailed data, an appropriate choice of moment—in this case, about $k \approx 3$—can yield substantial improvements. In particular, assumptions of data boundedness ($k = \infty$) give unnecessarily high variance and large radii r_k, while assuming too few moments ($k < 2$) yields a slow convergence rate $n^{-\frac{k-1}{k}}$ that causes absolute errors in estimation that are too large. In Table 1, for each privacy level $\alpha \in \{\frac{1}{10}, \frac{1}{2}, 1, 2, \infty\}$ (with $\alpha = \infty$ corresponding to no privacy) we tabulate the best mean absolute error achieved by any moment k (and the moment achieving this error). The table shows that, even performing a *post-hoc* optimal choice of moment estimator k, local differential privacy may be quite a strong constraint on the quality of estimation, as even with $\alpha = 1$ we incur an error of approximately 6% on average, while a non-private estimator observing half the population has error of .2%.

6.1.2 Median salary estimation

We now turn to an evaluation of the performance of the minimax optimal stochastic gradient procedure (11) for median estimation, comparing it with more naive but natural estimators based on noisy truncated versions of individuals' data. We first motivate the alternative strategy using the problem setting of Section 3.2.2. Recall that $\Pi_{[-r,r]}(x)$ denotes the projection of $x \in \mathbb{R}$ onto the interval $[-r, r]$. If by some prior knowledge we know that the distribution P satisfies $|\operatorname{med}(P)| \le r$, then for $X \sim P$ the random variables $\Pi_{[-r,r]}(X)$ have identical median to P, and for any symmetric random variable W, the variable $Z = W + \Pi_{[-r,r]}(X)$ satisfies $\operatorname{med}(Z) = \operatorname{med}(X)$. Now, let $W_i \stackrel{\text{iid}}{\sim} \text{Laplace}(\alpha/(2r))$, and consider the natural α-locally differentially private estimator

$$Z_i = \Pi_{[-r,r]}(X_i) + W_i, \quad \text{with} \quad \widehat{\theta}_n = \operatorname{med}(Z_1, \ldots, Z_n). \tag{42}$$

The variables Z_i are locally private versions of the X_i, as we simply add noise to (truncated) versions of the true data X_i. Rather than giving a careful theoretical investigation of the performance of the estimator (42), however, we turn to empirical results.

We again consider the salary data available from the UC system for fiscal year 2010, where we know that the median salary is at least zero (so we replace projections onto $[-r, r]$ with $[0, r]$ in the methods (11) and (42)). We compare the stochastic gradient estimator (11) using the specified step-sizes η_i and averaged predictor $\widehat{\theta}_n$ as well as the naive estimator (42). Treating the full population (all the salaries) as the sampling distribution P, we perform five experiments, each for a different estimate value of the radius of the median r. Specifically, we use $r \in \{1.5, 2, 4, 8, 16\} \operatorname{med}(P)$, where we assume that we know that the median must lie in some region near the true median. (The true median salary is approximately \$25,000.) For each experiment, we perform the following steps 400 times (with independent randomness) on the population of size $n = 252{,}540$:

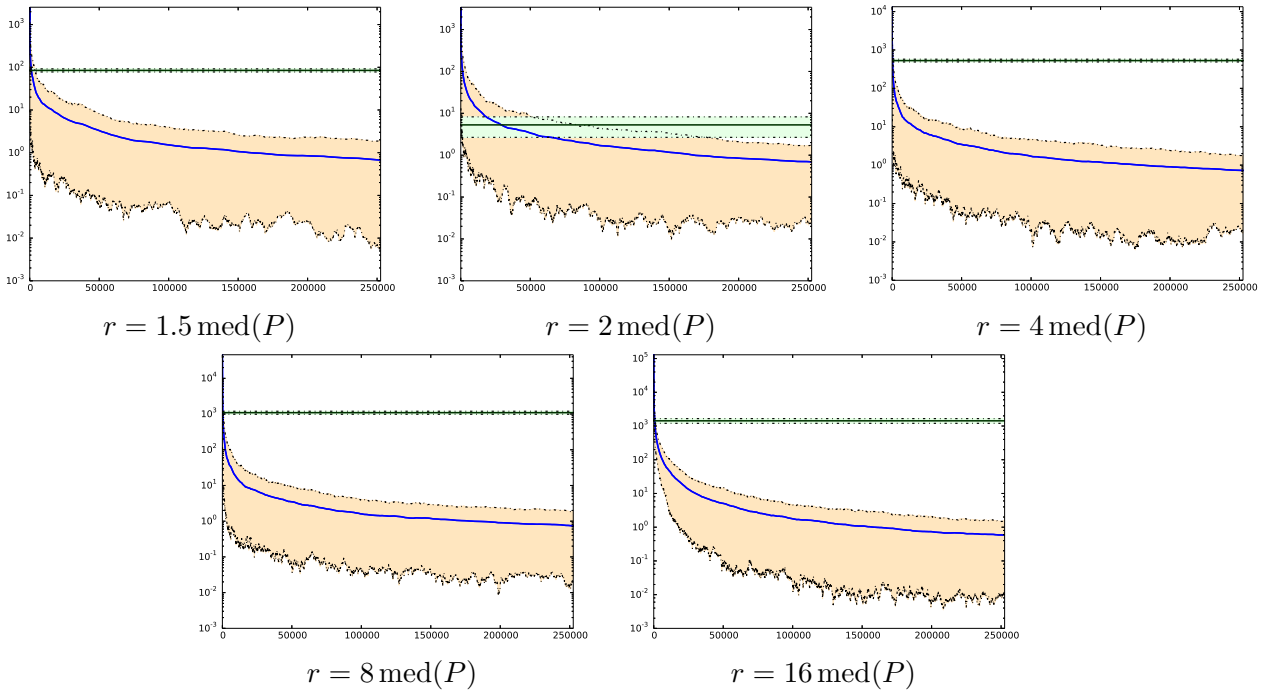

Figure 6. Performance of two median estimators—the SGD estimator (11) and the naive estimator (42)—plotted against sample size for $\alpha = 1$ for privately estimating median salary in the UC system. Each plot corresponds to a different estimated maximum radius r for the true median.

1. We perform n steps of private stochastic gradient descent (11), beginning from a uniformly random intialization $\theta_0 \in [0, r]$. Each step consists of a random draw X_i without replacement from the population $\{X_1, \ldots, X_n\}$, then performing the private gradient step.

2. We compute the naive private median (42) with i.i.d. $W_i \sim \mathrm{Laplace}(\alpha/(2r))$.

In Figure 6, we plot the results of these experiments. Each plot contains two lines. The first is a descending blue line (the lower of the two), which is the running gap $R(\widehat{\theta}_t) - R(\mathrm{med}(P))$ of the mean SGD estimator $\widehat{\theta}_t = \frac{1}{t}\sum_{i=1}^{t} \theta_i$ for each $t = 1, \ldots, n$; the gap is averaged over all 400 runs of SGD. We shade the region between the 5th and 95th percentile of the gaps $R(\widehat{\theta}_t) - R(\mathrm{med}(P))$ for the estimator over all 400 runs of SGD; note that the upward deviations are fairly tight. The flat line is the mean performance $R(\widehat{\theta}) - R(\mathrm{med}(P))$ of the naive estimator (42) over all 400 tests with 90% confidence bands above and below (this estimator uses the full sample of size n). Two aspects of these results are notable. First, the minimax optimal SGD estimator always has substantially better performance—even in the best case for the naive estimator (when $r = 2\,\mathrm{med}(P)$), the gap between the two is approximately a factor of 6—and usually gives a several order of magnitude improvement. Secondly, the optimal SGD estimator is quite robust, yielding mean performance $\mathbb{E}[R(\widehat{\theta}_n)] - R(\mathrm{med}(P)) < 1$ for all experiments. (By point of comparison, the non-private stochastic gradient estimator—that with $\alpha = +\infty$—attains error approximately .25.) More precisely, as our problem has median salary approximately \$25,000, the relative error in the gap—no matter which radius r is chosen—is at most $4 \cdot 10^{-5}$ for private SGD.

6.2 Drug use and hospital admissions

In this section, we consider estimation of the proportions of a population admitted to hospital emergency rooms for different types of drug use; it is natural that admitted persons might want to keep their drug use private, but accurate accounting of drug use is helpful for assessment of public health. We perform simulations based on real drug use data from the National Estimates of Drug-Related Emergency Department Visits (NEDREDV) [46] and treat this problem as a mean-estimation problem. First, we describe our data generation procedure, after which we describe our experiments in detail.

The NEDREDV data consists of tabulated triples of the form $(\text{drug}, \text{year}, m)$, where m is a count of the number of hospital admissions for patients using the given drug in the given year. We take admission data from the year 2004, which consists of 959,715 emergency department visits, and we include admissions for $d = 27$ common drugs.[2] Given these tuples, we generate a random dataset $\{X_1, \ldots, X_N\} \subset \mathbb{R}^d$ with $N = 959{,}715$ with the property that for each drug $j \in \{1, \ldots, d\}$, the marginal counts $\sum_{i=1}^N X_{ij}$ yield the correct drug use frequencies. Under this marginal constraint, we set each coordinate $X_{ij} = 0$ or 1 independently and uniformly at random. Thus, each non-private observation consists of a vector $X \in \{0,1\}^d$ representing a hospital admission, where coordinate j of X is 1 if the admittee abuses drug j and 0 otherwise. Many admittees abuse multiple drugs (this is true in the non-simulated data as there are substantially more drug counts than total admissions), so we consider the problem of estimating the mean $\frac{1}{N}\sum_{i=1}^N X_i$ of the population, $\theta = \mathbb{E}[X]$, where all we know is that $X \in [0,1]^d$.

In each separate experiment we draw a random sample of size $n = \frac{2N}{3}$ from the population, replacing each element X_i of the sample with an α-locally differentially private view Z_i, and construct an estimate $\widehat{\theta}$ of the true mean θ. In this case, the minimax optimal α-differentially private strategy is the ℓ_∞ sampling strategy (28). We also compare to a naive estimator that uses $Z = X + W$, where $W \in \mathbb{R}^d$ has independent coordinates each with Laplace(α/d) distribution, as well as a non-private estimator using the average of the subsampled vectors X.

Our results were shown earlier in Figure 1 in the Introduction, where we plot the results of 100 independent experiments with privacy parameter $\alpha = \frac{1}{2}$. We show the mean ℓ_∞ error for estimating the population proportions $\theta = \frac{1}{N}\sum_{i=1}^N X_i$—based on a population of size $N = 959715$—using a sample of size n as n ranges from 1 to $n = 600000$. We consider estimation of $d = 27$ drugs. The top-most (blue) line corresponds to the Laplace estimator, the bottom (black) line a non-private estimator based on empirical counts, and the middle (green) line the optimal private estimator. The plot also shows 5th and 95th percentile quantiles for each of the private experiments. From the figure, it is clear that the minimax optimal sampling strategy outperforms the equally private Laplace noise addition mechanism; even the worst performing random samples of the optimal sampling scheme outperform the best of the Laplace noise addition scheme. The mean error of the optimal scheme is also roughly a factor of $\sqrt{d} \approx 5$ better than the non-optimal Laplace scheme.

[2] The drugs are Alcohol, Cocaine, Heroin, Marijuana, Stimulants, Amphetamines, Methamphetamine, MDMA (Ecstasy), LSD, PCP, Antidepressants, Antipsychotics, Miscellaneous hallucinogens, Inhalants, lithium, Opiates, Opiates unspecified, Narcotic analgesics, Buprenorphine, Codeine, Fentanyl, Hydrocodone, Methadone, Morphine, Oxycodone, Ibuprofen, Muscle relaxants.

6.3 Censorship and privacy

In our final set of experiments, we investigate private estimation strategies for conditional probability estimation—specifically, logistic regression—in the context of prediction of whether a document will be censored. We have a collection of $N = 190{,}000$ Chinese blog posts, of which $N_{\text{cens}} = 90{,}000$ have been censored and $N_{\text{un}} = 100{,}000$ have been allowed to remain on Weibao (a Chinese blogging platform) by Chinese authorities.[3] The goal is to find those words strongly correlated with censorship decisions by estimation of conditional probabilities of censorship and word appearance. For each of the words appearing in 1.5% or more of the documents in the corpus (we estimate frequency on the entire corpus of $N = 190{,}000$ documents), we perform the following task. Let $X_i \in \{0, 1\}$ indicate whether the word appeared in document i, and let $Y_i \in \{-1, 1\}$ indicate whether document i was censored. For each word, our task is to estimate the conditional probabilities $P(Y = y \mid X = x)$ for $y \in \{-1, 1\}$ and $x \in \{0, 1\}$.

For this problem, we use the minimax optimal estimation procedure of Section 3.2.3, which consists of first estimating the joint probability $P(Y = y, X = x)$, treated as a multinomial parameter $p \in \Delta_4 = \{v \in \mathbb{R}^4_+ : \langle \mathbb{1}, v \rangle = 1\}$, then solving the logistic regression problem (13) with the estimated probability $\widehat{P}_n$. More concretely, we repeat the following experiment for each of the common words (those appearing in at least 1.5% of documents) a total of 50 times with privacy parameter $\alpha = 1$:

1. Draw a sample of $n = N/4$ of the documents, and construct an estimate $\widehat{P}_n$ of the joint probability $P(Y = y, X = x)$ using either
   (a) the randomized response multinomial estimator (34);
   (b) the multinomial estimator (34), except that the unbiased (unprojected) estimate of $\widehat{P}_n$ is constructed by adding independent Laplace(α) noise to each observed pair (X_i, Y_i);
   (c) the true counts (X_i, Y_i) in the sample of size n, which corresponds to standard (non-private) maximum likelihood estimation.
2. Construct $\widehat{\theta}_n$ by solving the logistic regression problem (13). We use radius $r = \log N$, as in the non-private problem it is impossible to achieve log-likelihood ratios of more than $\log N$.

We present results for these experiments in Figures 7 and 8. In each figure, we plot mean error for the 156 most common words (those appearing in at least 1.5% of documents) for the three estimators (1a)–(1c), that is, Laplace noise addition, randomized response, and no privacy. In the first figure, the horizontal axis gives word frequency and vertical axis is the average of $\|\widehat{\theta}_n - \theta\|_2$ computed over the 50 experiments. In this case, we construct θ by performing logistic regression on the entire dataset of $N = 190{,}000$ documents. While the plot is noisy, a few trends are clear. First, randomized response gives better performance than the Laplace noise addition technique for the same privacy level. For only one of the 156 words does Laplace noise addition yield better average performance for the error $\|\widehat{\theta}_n - \theta\|_2$. Letting $\widehat{\theta}^{\text{Lap}}_n$ denote the Laplace estimator and $\widehat{\theta}^{\text{rr}}_n$ the randomized-response-based estimator, the median ratio of $\|\widehat{\theta}^{\text{Lap}}_n - \theta\|_2 / \|\widehat{\theta}^{\text{rr}}_n - \theta\|_2$ across all words and experiments is 1.54, and the geometric mean of the ratio is 1.60. In addition, in Figure 7,

[3] We use data identical to that used in the articles [33, 34]. The datasets were constructed as follows: all blog posts from `http://weiboscope.jmsc.hku.hk/datazip/` were downloaded (see wa Fu et al. [48]) then segmented the Chinese text of each post using the Stanford Chinese language parser [37]. Of these, a random subsample of N_{cens} censored blog posts and N_{un} uncensored posts is taken.

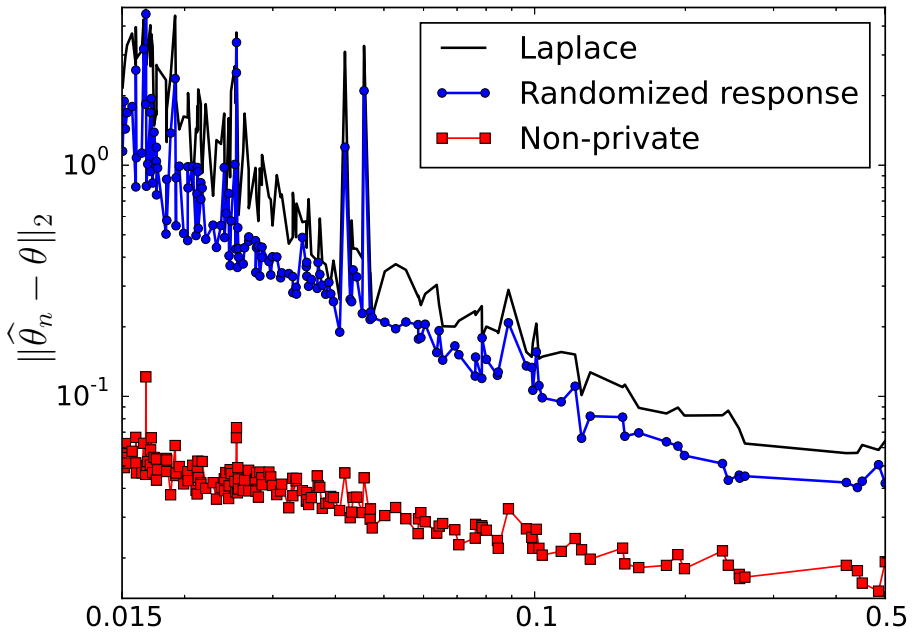

Figure 7: Mean ℓ_2-error for logistic regression, plotted against word frequency $P(X = 1)$.

the increase in mean error $\|\widehat{\theta}_n - \theta\|_2$ as the frequency $P(X = 1)$ decreases appears super-linear for the private estimators—worse than the non-private estimator—as predicted by our optimality results in Corollary 3. In Figure 8, we plot the average error $\|\widehat{\theta}_n - \theta\|_2$ for each of the estimation strategies along with standard error intervals in gray; here the word order (the permutation of entries for the horizontal axis) is chosen so that the mean error is decreasing for the randomized response estimator. Again, it is clear that the randomized response strategy dominates the more naive Laplace noise addition strategy, while non-private estimation enjoys improvements over both.

7 Conclusions

We have linked minimax analysis from statistical decision theory with differential privacy, bringing some of their respective foundational principles into close contact. Our main technique, in the form of the divergence inequalities in Theorems 1, 2, and 3, and their associated corollaries, shows that applying differentially private sampling schemes essentially acts as a contraction on distributions. These contractive inequalities allow us to provide results, presented in Propositions 1, 2, and 3, generalizing the classical minimax lower bounding techniques of Le Cam, Fano, and Assouad. These results allow us to give sharp minimax rates for estimation in locally private settings. With our examples in Sections 4.2, 5.2.1, and 5.2.2, we have developed a framework that shows that, roughly, if one can construct a family of distributions $\{P_\nu\}$ on the sample space $\mathcal{X}$ that is not well "correlated" with any member of $f \in L^\infty(\mathcal{X})$ for which $f(x) \in \{-1, 1\}$, then providing privacy is costly: the contraction provided in Theorems 2 and 3 is strong.

By providing sharp convergence rates for many standard statistical estimation procedures under local differential privacy, we have developed and explored some tools that may be used to better

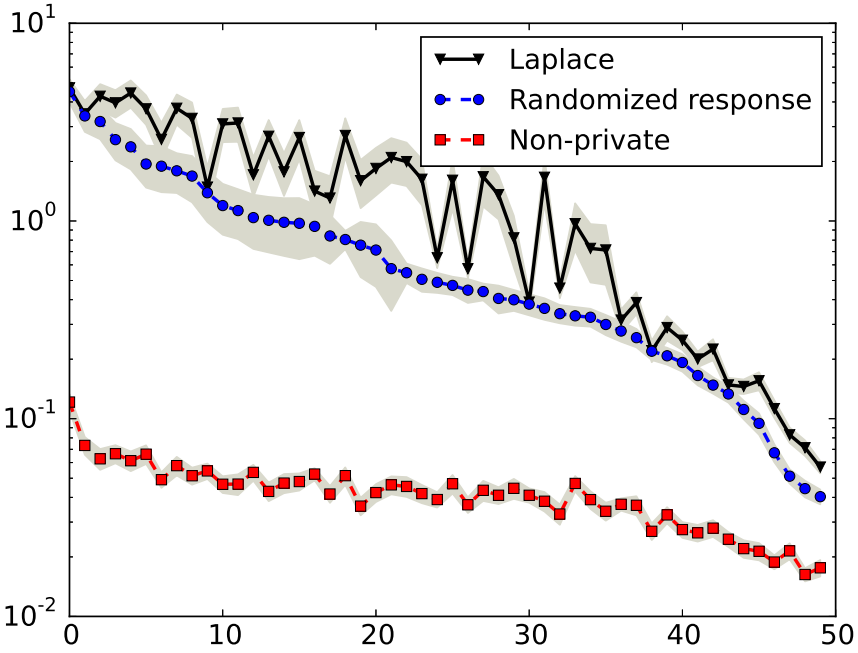

Figure 8. Mean ℓ_2-error for logistic regression, in sorted order of ℓ_2-error for randomized response.

understand privacy-preserving statistical inference. We have identified a fundamental continuum along which privacy may be traded for utility in the form of accurate statistical estimates, providing a way to adjust statistical procedures to meet the privacy or utility needs of the statistician and the population being sampled.

There are a number of open questions raised by our work. It is natural to wonder whether it is possible to obtain tensorized inequalities of the form of Corollary 4 even for interactive mechanisms. Another important question is whether the results we have provided can be extended to settings in which standard (non-local) differential privacy, or another form of disclosure limitation, holds. Such extensions could yield insights into optimal mechanisms for a number of private procedures.

Acknowledgments

We are very thankful to Shuheng Zhou for pointing out errors in Corollary 4 and Proposition 2 in an earlier version of this manuscript. We also thank Guy Rothblum for helpful discussions and Margaret Roberts for providing us tokenized data on the censorship task. Our work was supported in part by the U.S. Army Research Office under grant number W911NF-11-1-0391, Office of Naval Research MURI grant N00014-11-1-0688, and the National Science Foundation (NSF) grant CCF-1553086.

References

[1] A. Agarwal, P. L. Bartlett, P. Ravikumar, and M. J. Wainwright. Information-theoretic lower bounds on the oracle complexity of convex optimization. *IEEE Transactions on Information*

Theory, 58(5):3235–3249, 2012.

[2] V. Anantharam, A. Gohari, S. Kamath, and C. Nair. On maximal correlation, hypercontractivity, and the data processing inequality studied by Erkip and Cover. *arXiv:1304.6133 [cs.IT]*, 2013. URL http://arxiv.org/abs/1304.6133.

[3] E. Arias-Castro, E. Candés, and M. Davenport. On the fundamental limits of adaptive sensing. *IEEE Transactions on Information Theory*, 59(1):472–481, 2013.

[4] P. Assouad. Deux remarques sur l'estimation. *C. R. Academy Scientifique Paris Séries I Mathematics*, 296(23):1021–1024, 1983.

[5] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar. Privacy, accuracy, and consistency too: A holistic solution to contingency table release. In *Proceedings of the 26th ACM Symposium on Principles of Database Systems*, 2007.

[6] A. Beimel, K. Nissim, and E. Omri. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology*, volume 5157 of *Lecture Notes in Computer Science*, pages 451–468. Springer, 2008.

[7] A. Beimel, S. P. Kasiviswanathan, and K. Nissim. Bounds on the sample complexity for private learning and private data release. In *Proceedings of the 7th Theory of Cryptography Conference*, pages 437–454, 2010.

[8] A. Blum, K. Ligett, and A. Roth. A learning theory approach to non-interactive database privacy. In *Proceedings of the Fourtieth Annual ACM Symposium on the Theory of Computing*, 2008.

[9] P. Brucker. An $O(n)$ algorithm for quadratic knapsack problems. *Operations Research Letters*, 3(3):163–166, 1984.

[10] R. Carroll and P. Hall. Optimal rates of convergence for deconvolving a density. *Journal of the American Statistical Association*, 83(404):1184–1186, 1988.

[11] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.

[12] G. M. Clarke, C. A. Anderson, F. H. Pettersson, L. R. Cardon, A. P. Morris, and K. T. Zondervan. Basic statistical analysis in genetic case-control studies. *Nature Protocols*, 6(2): 121–133, 2011.

[13] T. M. Cover and J. A. Thomas. *Elements of Information Theory, Second Edition*. Wiley, 2006.

[14] J. C. Duchi and M. J. Wainwright. Distance-based and continuum Fano inequalities with applications to statistical estimation. *arXiv:1311.2669 [cs.IT]*, 2013.

[15] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. Privacy aware learning. *Journal of the Association for Computing Machinery*, 61(6), 2014.

[16] G. T. Duncan and D. Lambert. Disclosure-limited data dissemination. *Journal of the American Statistical Association*, 81(393):10–18, 1986.

[17] G. T. Duncan and D. Lambert. The risk of disclosure for microdata. *Journal of Business and Economic Statistics*, 7(2):207–217, 1989.

[18] C. Dwork and J. Lei. Differential privacy and robust statistics. In *Proceedings of the Fourty-First Annual ACM Symposium on the Theory of Computing*, 2009.

[19] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference*, pages 265–284, 2006.

[20] C. Dwork, G. N. Rothblum, and S. P. Vadhan. Boosting and differential privacy. In *51st Annual Symposium on Foundations of Computer Science*, pages 51–60, 2010.

[21] S. Efromovich. *Nonparametric Curve Estimation: Methods, Theory, and Applications*. Springer-Verlag, 1999.

[22] A. V. Evfimievski, J. Gehrke, and R. Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the Twenty-Second Symposium on Principles of Database Systems*, pages 211–222, 2003.

[23] S. E. Fienberg, U. E. Makov, and R. J. Steele. Disclosure limitation using perturbation and related methods for categorical data. *Journal of Official Statistics*, 14(4):485–502, 1998.

[24] S. E. Fienberg, A. Rinaldo, and X. Yang. Differential privacy and the risk-utility tradeoff for multi-dimensional contingency tables. In *The International Conference on Privacy in Statistical Databases*, 2010.

[25] S. R. Ganta, S. Kasiviswanathan, and A. Smith. Composition attacks and auxiliary information in data privacy. In *Proceedings of the 14th ACM SIGKDD Conference on Knowledge and Data Discovery (KDD)*, 2008.

[26] R. M. Gray. *Entropy and Information Theory*. Springer, 1990.

[27] R. Hall, A. Rinaldo, and L. Wasserman. Random differential privacy. *arXiv:1112.2680 [stat.ME]*, 2011. URL http://arxiv.org/abs/1112.2680.

[28] M. Hardt and G. N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *51st Annual Symposium on Foundations of Computer Science*, 2010.

[29] M. Hardt and K. Talwar. On the geometry of differential privacy. In *Proceedings of the Fourty-Second Annual ACM Symposium on the Theory of Computing*, pages 705–714, 2010. URL http://arxiv.org/abs/0907.3754.

[30] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, second edition, 2009.

[31] P. Kairouz, S. Oh, and P. Vishwanath. Extremal mechanisms for local differential privacy. In *Advances in Neural Information Processing Systems 27*, 2014.

[32] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

[33] G. King, J. Pan, and M. Roberts. How censorship in China allows government criticism but silences collective expression. *American Political Science Review*, 107(2):1–18, 2013.

[34] G. King, J. Pan, and M. Roberts. Reverse-engineering censorship in China: Randomized experimentation and participant observation. *Science*, 345(6199), 2014. doi: 10.1126/science.1251722.

[35] L. Le Cam. Convergence of estimates under dimensionality restrictions. *Annals of Statistics*, 1(1):38–53, 1973.

[36] J. Lei. Differentially private M-estimators. In *Advances in Neural Information Processing Systems 25*, 2011.

[37] R. Levy and C. D. Manning. Is it harder to parse Chinese, or the Chinese Treebank? In *Association for Computational Linguistics*, 2003.

[38] A. McGregor, I. Mironov, T. Pitassi, O. Reingold, K. Talwar, and S. Vadhan. The limits of two-party differential privacy. In *51st Annual Symposium on Foundations of Computer Science*, 2010.

[39] S. Negahban, P. Ravikumar, M. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, 27 (4):538–557, 2012.

[40] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

[41] U. of California. 2010 annual report on employee compensation, 2014. URL http://compensation.universityofcalifornia.edu/payroll2014/.

[42] R. R. Phelps. *Lectures on Choquet's Theorem, Second Edition*. Springer, 2001.

[43] B. I. P. Rubinstein, P. L. Bartlett, L. Huang, and N. Taft. Learning in a large function space: privacy-preserving mechanisms for SVM learning. *Journal of Privacy and Confidentiality*, 4 (1):65–100, 2012.

[44] D. Scott. On optimal and data-based histograms. *Biometrika*, 66(3):605–610, 1979.

[45] A. Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Fourty-Third Annual ACM Symposium on the Theory of Computing*, 2011.

[46] Substance Abuse and Mental Health Services Administration. Drug abuse warning network 2011: National estimates of drug-related emergency department visits. Technical Report SMA 13-4760, U.S. Department of Health and Human Services, 2013. URL http://www.samhsa.gov/data/emergency-department-data-dawn/reports?tab=47.

[47] A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009.

[48] K. wa Fu, C. Chan, and M. Chau. Assessing censorship on microblogs in china: Discriminatory keyword analysis and impact evaluation of the "real name registration" policy. *IEEE Internet Computing*, 17(3):42–50, 2013.

[49] S. Warner. Randomized response: a survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

[50] L. Wasserman and S. Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.

[51] Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, 27(5):1564–1599, 1999.

[52] B. Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer-Verlag, 1997.

A Proof of Theorem 1 and related results

We now collect proofs of our main results, beginning with Theorem 1.

A.1 Proof of Theorem 1

Observe that M_1 and M_2 are absolutely continuous with respect to one another, and there is a measure μ with respect to which they have densities m_1 and m_2, respectively. The channel probabilities $Q(\cdot \mid x)$ and $Q(\cdot \mid x')$ are likewise absolutely continuous, so that we may assume they have densities $q(\cdot \mid x)$ and write $m_i(z) = \int q(z \mid x) dP_i(x)$. In terms of these densities, we have

$$\begin{aligned} D_{\rm kl}\left(M_1 \| M_2\right) + D_{\rm kl}\left(M_2 \| M_1\right) &= \int m_1(z) \log \frac{m_1(z)}{m_2(z)} d\mu(z) + \int m_2(z) \log \frac{m_2(z)}{m_1(z)} d\mu(z) \\ &= \int \left(m_1(z) - m_2(z)\right) \log \frac{m_1(z)}{m_2(z)} d\mu(z). \end{aligned}$$

Consequently, we must bound both the difference $m_1 - m_2$ and the log ratio of the marginal densities. The following two auxiliary lemmas are useful:

Lemma 2. *For any α-locally differentially private channel Q, we have*

$$|m_1(z) - m_2(z)| \le c_\alpha \inf_x q(z \mid x) \left(e^\alpha - 1\right) \|P_1 - P_2\|_{\rm TV}, \tag{43}$$

where $c_\alpha = \min\{2, e^\alpha\}$.

Lemma 3. *Let $a, b \in \mathbb{R}_+$. Then $\left|\log \frac{a}{b}\right| \le \frac{|a-b|}{\min\{a,b\}}$.*

We prove these two results at the end of this section.

With the lemmas in hand, let us now complete the proof of the theorem. From Lemma 3, the log ratio is bounded as

$$\left|\log \frac{m_1(z)}{m_2(z)}\right| \le \frac{|m_1(z) - m_2(z)|}{\min\left\{m_1(z), m_2(z)\right\}}.$$

Applying Lemma 2 to the numerator yields

$$\begin{aligned} \left|\log \frac{m_1(z)}{m_2(z)}\right| &\le \frac{c_\alpha \left(e^\alpha - 1\right) \|P_1 - P_2\|_{\rm TV} \ \inf_x q(z \mid x)}{\min\{m_1(z), m_2(z)\}} \\ &\le \frac{c_\alpha \left(e^\alpha - 1\right) \|P_1 - P_2\|_{\rm TV} \ \inf_x q(z \mid x)}{\inf_x q(z \mid x)}, \end{aligned}$$

where the final step uses the inequality $\min\{m_1(z), m_2(z)\} \geq \inf_x q(z \mid x)$. Putting together the pieces leads to the bound

$$\left|\log \frac{m_1(z)}{m_2(z)}\right| \leq c_\alpha(e^\alpha - 1) \|P_1 - P_2\|_{\mathrm{TV}}.$$

Combining with inequality (43) yields

$$D_{\mathrm{kl}}(M_1 \| M_2) + D_{\mathrm{kl}}(M_2 \| M_1) \leq c_\alpha^2 (e^\alpha - 1)^2 \|P_1 - P_2\|_{\mathrm{TV}}^2 \int \inf_x q(z \mid x) d\mu(z).$$

The final integral is at most one, which completes the proof of the theorem.

It remains to prove Lemmas 2 and 3. We begin with the former. For any $z \in \mathcal{Z}$, we have

$$\begin{aligned}
m_1(z) - m_2(z) &= \int_{\mathcal{X}} q(z \mid x) [dP_1(x) - dP_2(x)] \\
&= \int_{\mathcal{X}} q(z \mid x) [dP_1(x) - dP_2(x)]_+ + \int_{\mathcal{X}} q(z \mid x) [dP_1(x) - dP_2(x)]_- \\
&\leq \sup_{x \in \mathcal{X}} q(z \mid x) \int_{\mathcal{X}} [dP_1(x) - dP_2(x)]_+ + \inf_{x \in \mathcal{X}} q(z \mid x) \int_{\mathcal{X}} [dP_1(x) - dP_2(x)]_- \\
&= \left(\sup_{x \in \mathcal{X}} q(z \mid x) - \inf_{x \in \mathcal{X}} q(z \mid x)\right) \int_{\mathcal{X}} [dP_1(x) - dP_2(x)]_+ .
\end{aligned}$$

By definition of the total variation norm, we have $\int [dP_1 - dP_2]_+ = \|P_1 - P_2\|_{\mathrm{TV}}$, and hence

$$|m_1(z) - m_2(z)| \leq \sup_{x, x'} \left|q(z \mid x) - q(z \mid x')\right| \|P_1 - P_2\|_{\mathrm{TV}}. \tag{44}$$

For any $\hat{x} \in \mathcal{X}$, we may add and subtract $q(z \mid \hat{x})$ from the quantity inside the supremum, which implies that

$$\begin{aligned}
\sup_{x, x'} \left|q(z \mid x) - q(z \mid x')\right| &= \inf_{\hat{x}} \sup_{x, x'} \left|q(z \mid x) - q(z \mid \hat{x}) + q(z \mid \hat{x}) - q(z \mid x')\right| \\
&\leq 2 \inf_{\hat{x}} \sup_{x} |q(z \mid x) - q(z \mid \hat{x})| \\
&= 2 \inf_{\hat{x}} q(z \mid \hat{x}) \sup_x \left|\frac{q(z \mid x)}{q(z \mid \hat{x})} - 1\right|.
\end{aligned}$$

Similarly, we have for any x, x'

$$|q(z \mid x) - q(z \mid x')| = q(z \mid x') \left|\frac{q(z \mid x)}{q(z \mid x')} - 1\right| \leq e^\alpha \inf_{\widehat{x}} q(z \mid \widehat{x}) \left|\frac{q(z \mid x)}{q(z \mid x')} - 1\right|.$$

Since for any choice of $x, \hat{x}$, we have $q(z \mid x)/q(z \mid \hat{x}) \in [e^{-\alpha}, e^\alpha]$, we find that (since $e^\alpha - 1 \geq 1 - e^{-\alpha}$)

$$\sup_{x, x'} \left|q(z \mid x) - q(z \mid x')\right| \leq \min\{2, e^\alpha\} \inf_x q(z \mid x) (e^\alpha - 1).$$

Combining with the earlier inequality (44) yields the claim (43).

To see Lemma 3, note that for any $x > 0$, the concavity of the logarithm implies that

$$\log(x) \le x - 1.$$

Setting alternatively $x = a/b$ and $x = b/a$, we obtain the inequalities

$$\log \frac{a}{b} \le \frac{a}{b} - 1 = \frac{a-b}{b} \quad \text{and} \quad \log \frac{b}{a} \le \frac{b}{a} - 1 = \frac{b-a}{a}.$$

Using the first inequality for $a \ge b$ and the second for $a < b$ completes the proof.

A.2 Proof of Corollary 4

Let us recall the definition of the induced marginal distribution (15), given by

$$M_\nu(S) = \int_{\mathcal{X}} Q(S \mid x_{1:n}) dP_\nu^n(x_{1:n}) \quad \text{for } S \in \sigma(\mathcal{Z}^n).$$

For each $i = 2, \ldots, n$, we let $M_{\nu(i)}(\cdot \mid Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1}) = M_{\nu(i)}(\cdot \mid z_{1:i-1})$ denote the (marginal over X_i) distribution of the variable Z_i conditioned on $Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1}$. In addition, use the shorthand notation

$$D_{\rm kl}\left(M_{\nu(i)} \| M_{\nu'(i)}\right) := \int_{\mathcal{Z}^{i-1}} D_{\rm kl}\left(M_{\nu(i)}(\cdot \mid z_{1:i-1}) \| M_{\nu'(i)}(\cdot \mid z_{1:i-1})\right) dM_\nu^{i-1}(z_1, \ldots, z_{i-1})$$

to denote the integrated KL divergence of the conditional distributions on the Z_i. By the chain rule for KL divergences [26, Chapter 5.3], we obtain

$$D_{\rm kl}\left(M_\nu^n \| M_{\nu'}^n\right) = \sum_{i=1}^n D_{\rm kl}\left(M_{\nu(i)} \| M_{\nu'(i)}\right).$$

By assumption (3), the distribution $Q_i(\cdot \mid X_i, Z_{1:i-1})$ for Z_i is α-differentially private for the sample X_i. As a consequence, if we let $P_{\nu(i)}(\cdot \mid Z_1 = z_1, \ldots, Z_{i-1} = z_{i-1})$ denote the conditional distribution of X_i given the first $i-1$ values $Z_1, \ldots, Z_{i-1}$ and the packing index $V = \nu$, then from the chain rule and Theorem 1 we obtain

$$\begin{aligned} D_{\rm kl}\left(M_\nu^n \| M_{\nu'}^n\right) &= \sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} D_{\rm kl}\left(M_{\nu(i)}(\cdot \mid z_{1:i-1}) \| M_{\nu'(i)}(\cdot \mid z_{1:i-1})\right) dM_\nu^{i-1}(z_{1:i-1}) \\ &\le \sum_{i=1}^n 4(e^\alpha - 1)^2 \int_{\mathcal{Z}^{i-1}} \left\| P_{\nu(i)}(\cdot \mid z_{1:i-1}) - P_{\nu'(i)}(\cdot \mid z_{1:i-1}) \right\|_{\rm TV}^2 dM_\nu^{i-1}(z_1, \ldots, z_{i-1}). \end{aligned}$$

By the construction of our sampling scheme, the random variables X_i are conditionally independent given $V = \nu$; thus the distribution $P_{\nu(i)}(\cdot \mid z_{1:i-1}) = P_{\nu(i)}$, where $P_{\nu(i)}$ denotes the distribution of X_i conditioned on $V = \nu$. Consequently, we have

$$\left\| P_{\nu(i)}(\cdot \mid z_{1:i-1}) - P_{\nu'(i)}(\cdot \mid z_{1:i-1}) \right\|_{\rm TV} = \left\| P_{\nu(i)} - P_{\nu'(i)} \right\|_{\rm TV},$$

which gives the claimed result.

B Proof of minimax bounds associated with Le Cam's method

In this appendix, we collect proofs of the various minimax lower bounds for specific problems in Section 3.

B.1 Proof of Corollary 1

The minimax rate characterized by equation (9) involves both a lower and an upper bound, and we divide our proof accordingly. We provide the proof for $\alpha \in (0, 1]$, but note that a similar result (modulo different constants) holds for any finite value of α.

Lower bound: We use Le Cam's method to prove the lower bound in equation (9). Fix a given constant $\delta \in (0, 1]$, with a precise value to be specified later. For $\nu \in \mathcal{V} \in \{-1, 1\}$, define the distribution P_ν with support $\{-\delta^{-1/k}, 0, \delta^{1/k}\}$ by

$$P_\nu(X = \delta^{-1/k}) = \frac{\delta(1+\nu)}{2}, \quad P_\nu(X = 0) = 1 - \delta, \quad \text{and} \quad P_\nu(X = -\delta^{-1/k}) = \frac{\delta(1-\nu)}{2}.$$

By construction, we have $\mathbb{E}[|X|^k] = \delta(\delta^{-1/k})^k = 1$ and $\theta_\nu = \mathbb{E}_\nu[X] = \delta^{\frac{k-1}{k}}\nu$, whence the mean difference is given by $\theta_1 - \theta_{-1} = 2\delta^{\frac{k-1}{k}}$. Applying Le Cam's method (7) yields

$$\mathfrak{M}_n(\Theta, (\cdot)^2, Q) \geq \left(\delta^{\frac{k-1}{k}}\right)^2 \left(\frac{1}{2} - \frac{1}{2}\left\|M_1^n - M_{-1}^n\right\|_{\mathrm{TV}}\right),$$

where M_ν^n denotes the marginal distribution of the samples $Z_1, \ldots, Z_n$ conditioned on $\theta = \theta_\nu$. Now Pinsker's inequality implies that $\left\|M_1^n - M_{-1}^n\right\|_{\mathrm{TV}}^2 \leq \frac{1}{2} D_{\mathrm{kl}}\left(M_1^n \| M_{-1}^n\right)$, and Corollary 4 yields

$$D_{\mathrm{kl}}\left(M_1^n \| M_{-1}^n\right) \leq 4(e^\alpha - 1)^2 n \left\|P_1 - P_{-1}\right\|_{\mathrm{TV}}^2 = 4(e^\alpha - 1)^2 n\delta^2.$$

Putting together the pieces yields $\left\|M_1^n - M_{-1}^n\right\|_{\mathrm{TV}} \leq (e^\alpha - 1)\delta\sqrt{2n}$. For $\alpha \in (0, 1]$, we have $e^\alpha - 1 \leq 2\alpha$, and thus our earlier application of Le Cam's method implies

$$\mathfrak{M}_n(\Theta, (\cdot)^2, \alpha) \geq \left(\delta^{\frac{k-1}{k}}\right)^2 \left(\frac{1}{2} - \alpha\delta\sqrt{2n}\right).$$

Substituting $\delta = \min\{1, 1/\sqrt{32n\alpha^2}\}$ yields the claim (9).

Upper bound: We must demonstrate an α-locally private conditional distribution Q and an estimator that achieves the upper bound in equation (9). We do so via a combination of truncation and addition of Laplacian noise. Define the truncation function $[\cdot]_T : \mathbb{R} \to [-T, T]$ by

$$[x]_T := \max\{-T, \min\{x, T\}\},$$

where the truncation level T is to be chosen. Let W_i be independent Laplace$(\alpha/(2T))$ random variables, and for each index $i = 1, \ldots, n$, define $Z_i := [X_i]_T + W_i$. By construction, the random variable Z_i is α-differentially private for X_i. For the mean estimator $\widehat{\theta} := \frac{1}{n}\sum_{i=1}^n Z_i$, we have

$$\mathbb{E}\left[(\widehat{\theta} - \theta)^2\right] = \mathrm{Var}(\widehat{\theta}) + \left(\mathbb{E}[\widehat{\theta}] - \theta\right)^2 = \frac{4T^2}{n\alpha^2} + \frac{1}{n}\mathrm{Var}([X_1]_T) + \left(\mathbb{E}[Z_1] - \theta\right)^2. \tag{45}$$

We claim that

$$\mathbb{E}[Z] = \mathbb{E}\left[[X]_T\right] \in \left[\mathbb{E}[X] - \frac{2^{1/k}}{T^{k-1}}, \mathbb{E}[X] + \frac{2^{1/k}}{T^{k-1}}\right]. \tag{46}$$

Indeed, by the assumption that $\mathbb{E}[|X|^k] \le 1$, if we take $k' = \frac{k}{k-1}$ so that $1/k' + 1/k = 1$, we have that

$$\begin{aligned}
\mathbb{E}[|X - [X]_T|] &= \mathbb{E}[|X - T|\mathbf{1}\{X > T\}] + \mathbb{E}[|X + T|\mathbf{1}\{X < -T\}] \\
&\le \mathbb{E}[|X|\mathbf{1}\{X > T\}] + \mathbb{E}[|X|\mathbf{1}\{X < -T\}] \\
&\le \mathbb{E}[|X|^k]^{1/k}\mathbb{P}(X > T)^{1/k'} + \mathbb{E}[|X|^k]^{1/k}\mathbb{P}(X < -T)^{1/k'} \le 2^{1/k}T^{1-k}.
\end{aligned}$$

The final inequality follows from Markov's inequality, as $\mathbb{P}(X > T) + \mathbb{P}(X < -T) = \mathbb{P}(|X| > T) \le T^{-k}$, and for any $k' \in [1, \infty]$, we have $\sup_{a+b \le c}\{a^{1/k'} + b^{1/k'} : a, b \ge 0\} = 2^{1-1/k'}c^{1/k'} = 2^{1/k}c^{1/k'}$.

From the bound (45) and the inequalities that since $[X]_T \in [-T, T]$ and $\alpha^2 \le 1$, we have

$$\mathbb{E}\left[(\widehat{\theta} - \theta)^2\right] \le \frac{5T^2}{n\alpha^2} + \frac{4}{T^{2k-2}} \quad \text{valid for any } T > 0.$$

Choosing $T = (n\alpha^2)^{1/(2k)}$ yields the upper bound (9) immediately.

B.2 Proof of Corollary 2

We have already given the proof of the upper bound in the proposition in the course of our discussion of the stochastic gradient descent estimator (11). We thus focus on the lower bound. Fix $r > 0$ and let $\delta \in [0, 1]$ be a constant to be chosen subsequently. For $\nu \in \{-1, 1\}$ consider the distributions P_ν supported on $\{-r, r\}$ defined by

$$P_\nu(X = r) = \frac{1 + \delta\nu}{2} \quad \text{and} \quad P_\nu(X = -r) = \frac{1 - \delta\nu}{2}.$$

We have $\text{med}(P_1) = r$ and $\text{med}(P_{-1}) = -r$, and using an extension of the classical reduction from minimax estimation to testing to gaps in function values (recall inequality (7); see also [1]), we claim that

$$\mathfrak{M}_n([-r, r], R, Q) \ge \frac{r\delta}{2} \inf_\psi \mathbb{P}\left(\psi(Z_1, \ldots, Z_n) \ne V\right), \tag{47}$$

where $V \in \{-1, 1\}$ is chosen at random and conditional on $V = \nu$, we draw $Z_i \stackrel{\text{iid}}{\sim} M_\nu$ for $M_\nu(\cdot) = \int Q(\cdot \mid x) dP_\nu(x)$. Indeed, under this model, by using the shorthand $R_\nu(\theta) = \mathbb{E}_{P_\nu}[|\theta - X|]$, we have $\inf_\theta R_\nu(\theta) = \frac{1-\delta}{2}r$; and for any estimator $\widehat{\theta}$ we have

$$\frac{1}{2}\sum_{\nu \in \{\pm 1\}} \left(\mathbb{E}_{P_\nu}[R_\nu(\widehat{\theta})] - \frac{1-\delta}{2}r\right) \ge \frac{1}{2}\sum_{\nu \in \{\pm 1\}} \mathbb{E}_{P_\nu}\left[\mathbf{1}\left\{\text{sign}(\widehat{\theta}) \ne \nu\right\}\frac{r\delta}{2}\right] \ge \frac{r\delta}{2}\inf_\psi \mathbb{P}(\psi(Z_{1:n}) \ne V)$$

as claimed. Continuing, we use Pinsker's inequality and Corollary 4, which gives

$$\left\|M_1^n - M_{-1}^n\right\|_{\text{TV}}^2 \le \frac{1}{4}\left[D_{\text{kl}}\left(M_1^n \| M_{-1}^n\right) + D_{\text{kl}}\left(M_{-1}^n \| M_1^n\right)\right] \le (e^\alpha - 1)^2 n \left\|P_1 - P_{-1}\right\|_{\text{TV}}^2 \le 3n\alpha^2\delta^2$$

for $\alpha \le 1$. Thus, the bound (47) implies that for any $\delta \in [0, 1]$ and any α-private channel Q, we have

$$\mathfrak{M}_n([-r, r], R, Q) \ge \frac{r\delta}{4}\left(1 - 3n\alpha^2\delta^2\right).$$

Take $\delta^2 = \frac{1}{6n\alpha^2}$ to give the result of the corollary with $c_\ell = \frac{1}{20}$.

B.3 Proof of Corollary 3

Lower bound We first prove the lower bound (12a). We use the α-private version (8) of Le Cam's inequality from Proposition 1 to prove the result. We prove the result under the assumption that $\mu \le \frac{1}{2}$; we show subsequently this is no loss of generality. When $\mu \le \frac{1}{2}$, we may assume that $\theta_0 = 0$ and prove a minimax lower bound in this setting. Fix $\delta \in [0, r]$, and consider the two conditional distributions

$$P_0(Y = y \mid X = x) = \frac{1}{1 + \exp(-y(r - \delta)x)} \quad \text{and} \quad P_1(Y = y \mid X = x) = \frac{1}{1 + \exp(-yrx)}.$$

Then for $x = 0$ we have $P_0(\cdot \mid X = x) = P_1(\cdot \mid X = x)$, while for $x = 1$ we have

$$\begin{aligned}\max_y |P_0(Y = y \mid X = 1) - P_1(Y = y \mid X = 1)| &= \max\left\{\left|\frac{1}{1 + e^{r-\delta}} - \frac{1}{1 + e^r}\right|, \left|\frac{1}{1 + e^{\delta - r}} - \frac{1}{1 + e^{-r}}\right|\right\} \\ &= \frac{e^r - e^{r-\delta}}{(1 + e^{r-\delta})(1 + e^r)} \le \frac{e^r}{e^{2r-\delta}}(1 - e^{-\delta}) = e^{-r}(e^\delta - 1).\end{aligned}$$

Recalling that $\mu = P(X = 1)$ and $\{P_\nu^{\text{joint}}\}_{\nu \in \{0,1\}}$, denotes the joint distribution of (X, Y), we obtain

$$\left\|P_0^{\text{joint}} - P_1^{\text{joint}}\right\|_{\text{TV}} = \mu \max_y |P_0(Y = y \mid X = 1) - P_1(Y = y \mid X = 1)| \le \mu e^{-r}(e^\delta - 1).$$

Applying the inequality (8), we find that no matter the distribution of X, we have

$$\mathfrak{M}_n(\theta(\mathcal{P}_{r,\mu}), |\cdot|^2, \alpha) \ge \frac{\delta}{4}\left(1 - 2\sqrt{n\alpha^2\mu^2} \cdot e^{-r}(e^\delta - 1)\right).$$

Using the fact that

$$2\alpha\sqrt{n\mu^2}e^{-r}(e^\delta - 1) \le \frac{1}{2} \quad \text{iff} \quad e^\delta \le 1 + \frac{e^r}{4\alpha\sqrt{n\mu^2}} \quad \text{or} \quad \delta \le \log\left(1 + \frac{e^r}{4\alpha\sqrt{n\mu^2}}\right)$$

gives the desired result by taking $\delta = \min\{\log(1 + e^r/4\alpha\mu\sqrt{n}), r\}$. If initially we had that $\mu > \frac{1}{2}$, then using the two conditional distributions

$$P_0(Y = y \mid X = x) = \frac{1}{1 + \exp(-y[(r - \delta) + \delta x])} \quad \text{and} \quad P_1(Y = y \mid X = x) = \frac{1}{1 + \exp(-yr)}$$

allows exactly the same derivation of a lower bound, but μ is replaced by $1 - \mu$ throughout the derivation.

Upper bound We now turn to the upper bound (12b). We first prove a distributional result, then show how this is sufficient to give the minimax guarantee. For shorthand, let us write the joint distribution $p_{11} = P(Y = 1, X = 1)$, $p_{-11} = P(Y = -1, X = 1)$, $p_{10} = P(Y = 1, X = 0)$, and $p_{-10} = P(Y = -1, X = 0)$. Then the true negative log likelihood is

$$\mathbb{E}_P[\log(1+e^{-Y(\theta_0+\theta_1 X)})] = p_{11}\log(1+e^{-\theta_0-\theta_1})+p_{-11}\log(1+e^{\theta_0+\theta_1})+p_{10}\log(1+e^{-\theta_0})+p_{-10}\log(1+e^{\theta_0}),$$

which is (by inspection) minimized by $\theta_0 = \log \frac{p_{10}}{p_{-10}}$ and $\theta_1 = \log \frac{p_{11}}{p_{-11}} - \theta_0$. We now show how the estimator (13) satisfies asymptotic normality guarantees. Recall the construction of $U(x, y) \in \{e_1, \ldots, e_4\}$ as

$$U(x,y) = \{e_1 \text{ if } y = 1, x = 0, \quad e_2 \text{ if } y = -1, x = 0, \quad e_3 \text{ if } y = 1, x = 1, \quad e_4 \text{ if } y = -1, x = 1\},$$

and let $p = \mathbb{E}[U] \in \Delta_4$, so that p corresponds to the multinomial joint probability vector of $P(Y = y, X = x)$. Let D be a diagonal matrix with independent entries, where $D_{jj} = 1$ with probability $q_\alpha = e^{\alpha/2}/(1 + e^{\alpha/2})$ and $D_{jj} = 0$ otherwise; let $D^{(i)}$, $i = 1, \ldots, n$, be an i.i.d. sequence of these matrices. Then the multinomial randomized response strategy (33) gives

$$Z_i = D^{(i)} U_i + (I - D^{(i)})(\mathbb{1} - U_i) \in \{0,1\}^4, \quad \text{which satisfies } \mathbb{E}[Z_i] = \frac{e^\alpha - 1}{e^\alpha + 1} p + \frac{1}{1 + e^\alpha} \mathbb{1}.$$

Recalling that $q = q_\alpha$ is the Bernoulli probability of flipping a response in our randomized response model, then a somewhat lengthy but straightforward algebraic calculation reveals that

$$\mathrm{Cov}(Z_i) = \mathrm{Cov}(DU + (I - D)(\mathbb{1} - U)) = (2q - 1)^2 (\mathrm{diag}(p) - pp^\top) + q(1 - q) I_{4\times 4}. \tag{48}$$

As a consequence of the covariance calcuation (48), if we define $\Sigma_{p,\alpha} = \mathrm{Cov}(Z_i)$ and let

$$\widehat{p}_n^0 := \frac{e^\alpha + 1}{e^\alpha - 1} \left(\frac{1}{n} \sum_{i=1}^n \left[Z_i - \frac{1}{1 + e^\alpha} \mathbb{1} \right] \right) = \frac{1}{2q - 1} \left(\frac{1}{n} \sum_{i=1}^n [Z_i - (1 - q)\mathbb{1}] \right),$$

we have $\mathbb{E}[\widehat{p}_n^0] = p$ and the convergence guarantee

$$\sqrt{n}(\widehat{p}_n^0 - p) \overset{d}{\rightsquigarrow} \mathsf{N}\left(0, \frac{1}{(2q-1)^2} \Sigma_{p,\alpha}\right),$$

where

$$(I - \frac{1}{4}\mathbb{1}\mathbb{1}^\top)\Sigma_{p,\alpha}(I - \frac{1}{4}\mathbb{1}\mathbb{1}^\top) = (2q-1)^2(\mathrm{diag}(p) - pp^\top) + q(1-q)(I - \frac{1}{4}\mathbb{1}\mathbb{1}^\top).$$

Now we use the delta method to show that our proposed estimator satisfies the claimed asymptotic guarantee. As $p \succ 0$ and $\widehat{p}_n^0 \overset{a.s.}{\to} p$, we have $\widehat{p}_n^0 \succ 0$ strictly eventually almost surely. As a consequence, the projection of $\widehat{p}_n^0$ onto the probability simplex is eventually given by the projection of an arbitrary vector onto the hyperplane $\{u \in \mathbb{R}^4 : \langle \mathbb{1}, u \rangle = 1\}$; that is, the mapping $p \mapsto (I - \mathbb{1}\mathbb{1}^\top/4)p + \mathbb{1}/4$, so that if $\widehat{p}_n = \Pi_{\Delta_4}(\widehat{p}_n^0)$, we have

$$\sqrt{n}\,(\widehat{p}_n - p) \overset{d}{\rightsquigarrow} \mathsf{N}\left(0, (2q-1)^{-2}(I - \frac{1}{4}\mathbb{1}\mathbb{1}^\top)\Sigma_{p,\alpha}(I - \frac{1}{4}\mathbb{1}\mathbb{1}^\top)\right).$$

Now we apply the delta method to the sequence $\widehat{\theta}_n - \theta$. The Jacobian of the mapping $p \mapsto \theta(p)$ is by inspection

$$J_\theta(p) := \theta'(p) = \begin{bmatrix} \frac{1}{p_{10}} & -\frac{1}{p_{-10}} & 0 & 0 \\ -\frac{1}{p_{10}} & \frac{1}{p_{-10}} & \frac{1}{p_{11}} & \frac{-1}{p_{-11}} \end{bmatrix} \in \mathbb{R}^{2\times 4},$$

so

$$\sqrt{n}(\widehat{\theta}_n - \theta) \overset{d}{\rightsquigarrow} \mathsf{N}\left(0, (2q-1)^{-2} J_\theta(p)(I - \frac{1}{4}\mathbb{1}\mathbb{1}^\top)\Sigma_{p,\alpha}(I - \frac{1}{4}\mathbb{1}\mathbb{1}^\top) J_\theta(p)^\top\right). \tag{49}$$

To simplify somewhat, we note that

$$J_\theta(p)p = 0, \quad J_\theta(p)\operatorname{diag}(p)J_\theta(p)^\top = \sum_y \begin{bmatrix} \frac{1}{p_{y0}} & -\frac{1}{p_{y0}} \\ -\frac{1}{p_{y0}} & \sum_x \frac{1}{p_{yx}} \end{bmatrix}, \quad J_\theta(p)J_\theta(p)^\top = \sum_y \begin{bmatrix} \frac{1}{p_{y0}^2} & -\frac{1}{p_{y0}^2} \\ -\frac{1}{p_{y0}^2} & \sum_x \frac{1}{p_{yx}^2} \end{bmatrix},$$

and

$$J_\theta(p)\mathbb{1}\mathbb{1}^\top J_\theta(p)^\top = \begin{bmatrix} (\frac{1}{p_{10}} - \frac{1}{p_{-10}})^2 & (\frac{1}{p_{10}} - \frac{1}{p_{-10}})(\frac{1}{p_{-10}} - \frac{1}{p_{10}} + \frac{1}{p_{11}} - \frac{1}{p_{-11}}) \\ (\frac{1}{p_{10}} - \frac{1}{p_{-10}})(\frac{1}{p_{-10}} - \frac{1}{p_{10}} + \frac{1}{p_{11}} - \frac{1}{p_{-11}}) & (\frac{1}{p_{-10}} - \frac{1}{p_{10}} + \frac{1}{p_{11}} - \frac{1}{p_{-11}})^2 \end{bmatrix}.$$

Taking the trace of the final covariance in the convergence guarantee (49), we have

$$\begin{aligned}
&\operatorname{tr}(J_\theta(p)(I - (1/4)\mathbb{1}\mathbb{1}^\top)\Sigma_{p,\alpha}(I - (1/4)\mathbb{1}\mathbb{1}^\top)J_\theta(p)^\top) \\
&= (2q-1)^2 \operatorname{tr}(J_\theta(p)\operatorname{diag}(p)J_\theta(p)^\top) + q(1-q)\operatorname{tr}(J_\theta(p)(I - (1/4)\mathbb{1}\mathbb{1}^\top)J_\theta(p)^\top) \\
&= (2q-1)^2 \sum_y \left(\frac{2}{p_{y0}} + \frac{1}{p_{y1}}\right) + q(1-q)\Bigg[\sum_y \left(\frac{2}{p_{y0}^2} + \frac{1}{p_{y1}^2}\right) \quad \cdots \\
&\qquad\qquad - \frac{1}{4}\left(\left(p_{10}^{-1} - p_{-10}^{-1}\right)^2 + (p_{-10}^{-1} - p_{10}^{-1} + p_{11}^{-1} - p_{-11}^{-1})^2\right)\Bigg].
\end{aligned}$$

This gives the desired result after we divide both terms in the summation by $(2q-1)^2$ and set $C_{\alpha,\theta}^2 = (2q-1)^{-2}\operatorname{tr}(J_\theta(p)(I - (1/4)\mathbb{1}\mathbb{1}^\top)\Sigma_{p,\alpha}(I - (1/4)\mathbb{1}\mathbb{1}^\top)J_\theta(p)^\top)$.

C Proof of Theorem 2 and related results

In this section, we collect together the proof of Theorem 2 and related corollaries.

C.1 Proof of Theorem 2

Let $\mathcal{Z}$ denote the domain of the random variable Z. We begin by reducing the problem to the case when $\mathcal{Z} = \{1, 2, \ldots, k\}$ for an arbitrary positive integer k. Indeed, in the general setting, we let $\mathcal{K} = \{K_i\}_{i=1}^k$ be any (measurable) finite partition of $\mathcal{Z}$, where for $z \in \mathcal{Z}$ we let $[z]_\mathcal{K} = K_i$ for the K_i such that $z \in K_i$. The KL divergence $D_{\rm kl}\left(M_\nu \| \overline{M}\right)$ can be defined as the supremum of the (discrete) KL divergences between the random variables $[Z]_\mathcal{K}$ sampled according to M_ν and $\overline{M}$ over all partitions $\mathcal{K}$ of $\mathcal{Z}$; for instance, see Gray [26, Chapter 5]. Consequently, we can prove the claim for $\mathcal{Z} = \{1, 2, \ldots, k\}$, and then take the supremum over k to recover the general case. Accordingly, we can work with the probability mass functions $m(z \mid \nu) = M_\nu(Z = z)$ and $\overline{m}(z) = \overline{M}(Z = z)$, and we may write

$$D_{\rm kl}\left(M_\nu \| \overline{M}\right) + D_{\rm kl}\left(\overline{M} \| M_\nu\right) = \sum_{z=1}^k (m(z \mid \nu) - \overline{m}(z)) \log \frac{m(z \mid \nu)}{\overline{m}(z)}. \tag{50}$$

Throughout, we will also use (without loss of generality) the probability mass functions $q(z \mid x) = Q(Z = z \mid X = x)$, where we note that $m(z \mid \nu) = \int q(z \mid x) dP_\nu(x)$.

Now we use Lemma 3 from the proof of Theorem 1 to complete the proof of Theorem 2. Starting with equality (50), we have

$$\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\left[D_{\rm kl}\left(M_\nu\|\overline{M}\right)+D_{\rm kl}\left(\overline{M}\|M_\nu\right)\right] \le \sum_{\nu\in\mathcal{V}}\frac{1}{|\mathcal{V}|}\sum_{z=1}^{k}|m(z\mid\nu)-\overline{m}(z)|\left|\log\frac{m(z\mid\nu)}{\overline{m}(z)}\right| \le \sum_{\nu\in\mathcal{V}}\frac{1}{|\mathcal{V}|}\sum_{z=1}^{k}|m(z\mid\nu)-\overline{m}(z)|\,\frac{|m(z\mid\nu)-\overline{m}(z)|}{\min\{\overline{m}(z),m(z\mid\nu)\}}.$$

Now, we define the measure m^0 on $\mathcal{Z}=\{1,\ldots,k\}$ by $m^0(z):=\inf_{x\in\mathcal{X}}q(z\mid x)$. It is clear that $\min\{\overline{m}(z),m(z\mid\nu)\}\ge m^0(z)$, whence we find

$$\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\left[D_{\rm kl}\left(M_\nu\|\overline{M}\right)+D_{\rm kl}\left(\overline{M}\|M_\nu\right)\right]\le\sum_{\nu\in\mathcal{V}}\frac{1}{|\mathcal{V}|}\sum_{z=1}^{k}\frac{(m(z\mid\nu)-\overline{m}(z))^2}{m^0(z)}.$$

It remains to bound the final sum. For any constant $c\in\mathbb{R}$, we have

$$m(z\mid\nu)-\overline{m}(z)=\int_{\mathcal{X}}(q(z\mid x)-c)\left(dP_\nu(x)-d\overline{P}(x)\right).$$

We define a set of functions $f:\mathcal{Z}\times\mathcal{X}\to\mathbb{R}$ (depending implicitly on q) by

$$\mathcal{F}_\alpha:=\left\{f\mid f(z,x)\in[1,e^\alpha]m^0(z)\text{ for all }z\in\mathcal{Z}\text{ and }x\in\mathcal{X}\right\}.$$

By the definition of differential privacy, when viewed as a joint mapping from $\mathcal{Z}\times\mathcal{X}\to\mathbb{R}$, the conditional p.m.f. q satisfies $\{(z,x)\mapsto q(z\mid x)\}\in\mathcal{F}_\alpha$. Since constant (with respect to x) shifts do not change the above integral, we can modify the range of functions in $\mathcal{F}_\alpha$ by subtracting $m^0(z)$ from each, yielding the set

$$\mathcal{F}'_\alpha:=\left\{f\mid f(z,x)\in[0,e^\alpha-1]\,m^0(z)\text{ for all }z\in\mathcal{Z}\text{ and }x\in\mathcal{X}\right\}.$$

As a consequence, we find that

$$\sum_{\nu\in\mathcal{V}}(m(z\mid\nu)-\overline{m}(z))^2\le\sup_{f\in\mathcal{F}_\alpha}\left\{\sum_{\nu\in\mathcal{V}}\left(\int_{\mathcal{X}}f(z,x)\left(dP_\nu(x)-d\overline{P}(x)\right)\right)^2\right\} =\sup_{f\in\mathcal{F}'_\alpha}\left\{\sum_{\nu\in\mathcal{V}}\left(\int_{\mathcal{X}}\left(f(z,x)-m^0(z)\right)\left(dP_\nu(x)-d\overline{P}(x)\right)\right)^2\right\}.$$

By inspection, when we divide by $m^0(z)$ and recall the definition of the set $\mathbb{B}_\infty\subset L^\infty(\mathcal{X})$ in the statement of Theorem 2, we obtain

$$\sum_{\nu\in\mathcal{V}}(m(z\mid\nu)-\overline{m}(z))^2\le\left(m^0(z)\right)^2(e^\alpha-1)^2\sup_{\gamma\in\mathbb{B}_\infty}\sum_{\nu\in\mathcal{V}}\left(\int_{\mathcal{X}}\gamma(x)\left(dP_\nu(x)-d\overline{P}(x)\right)\right)^2.$$

Putting together our bounds, we have

$$\begin{aligned}
&\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\left[D_{\rm kl}\left(M_\nu\|\overline{M}\right)+D_{\rm kl}\left(\overline{M}\|M_\nu\right)\right] \\
&\quad\le (e^\alpha-1)^2\sum_{z=1}^{k}\frac{1}{|\mathcal{V}|}\frac{\left(m^0(z)\right)^2}{m^0(z)}\sup_{\gamma\in\mathbb{B}_\infty}\sum_{\nu\in\mathcal{V}}\left(\int_{\mathcal{X}}\gamma(x)\left(dP_\nu(x)-d\overline{P}(x)\right)\right)^2 \\
&\quad\le (e^\alpha-1)^2\frac{1}{|\mathcal{V}|}\sup_{\gamma\in\mathbb{B}_\infty}\sum_{\nu\in\mathcal{V}}\left(\int_{\mathcal{X}}\gamma(x)\left(dP_\nu(x)-d\overline{P}(x)\right)\right)^2,
\end{aligned}$$

since $\sum_z m^0(z)\le 1$, which is the statement of the theorem.

C.2 Proof of Inequality (24)

In the non-interactive setting (4), the marginal distribution M_ν^n is a product measure and Z_i is conditionally independent of $Z_{1:i-1}$ given V. Thus by the chain rule for mutual information [26, Chapter 5] and the fact (as in the proof of Theorem 2) that we may assume w.l.o.g. that Z has finite range

$$I(Z_1,\ldots,Z_n;V)=\sum_{i=1}^{n}I(Z_i;V\mid Z_{1:i-1})=\sum_{i=1}^{n}\left[H(Z_i\mid Z_{1:i-1})-H(Z_i\mid V,Z_{1:i-1})\right].$$

Since conditioning reduces entropy and $Z_{1:i-1}$ is conditionally independent of Z_i given V, we have $H(Z_i\mid Z_{1:i-1})\le H(Z_i)$ and $H(Z_i\mid V,Z_{1:i-1})=H(Z_i\mid V)$. In particular, we have

$$I(Z_1,\ldots,Z_n;V)\le\sum_{i=1}^{n}I(Z_i;V)=\sum_{i=1}^{n}\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}D_{\rm kl}\left(M_{\nu,i}\|\overline{M}_i\right).$$

Applying Theorem 2 completes the proof.

D Proofs of multi-dimensional mean-estimation results

At a high level, our proofs of these results consist of three steps, the first of which is relatively standard, while the second two exploit specific aspects of the local privacy setting. We outline them here:

(1) The first step is an essentially standard reduction, based on inequality (21) in Section 4, from an estimation problem to a multi-way testing problem that involves discriminating between indices ν contained within some subset $\mathcal{V}$ of $\mathbb{R}^d$.

(2) The second step is an appropriate construction of a packing set $\mathcal{V}\subset\mathbb{R}^d$. We require the existence of a well-separated set: one for which ratio of the packing set size $|\mathcal{V}|$ to neighborhood size N_δ is large enough relative to the separation δ of definition (19).

(3) The final step is to apply Theorem 2 in order to control the mutual information associated with the testing problem. Doing so requires bounding the supremum in Theorem 2 (and inequality (24)) via the the operator norm of $\mathrm{Cov}(V)$ for a vector V drawn uniformly at random from $\mathcal{V}$. This is made easier by the uniformity of the sampling scheme allowed by the generalization (21) of Fano's inequality we use, as it is possible to enforce that $\mathrm{Cov}(V)$ has relatively small operator norm.

The estimation to testing reduction of Step 1 is accomplished by the reduction (21) of Section 4. Accordingly, the proofs to follow are devoted to the second and third steps in each case.

D.1 Proof of Corollary 5

We provide a proof of the lower bound, as we provided the argument for the upper bound in Section 4.2.3.

Constructing a well-separated set: Let k be an arbitrary integer in $\{1, 2, \ldots, d\}$, and let $\mathcal{V}_k = \{-1, 1\}^k$ denote the k-dimensional hypercube. We extend the set $\mathcal{V}_k \subseteq \mathbb{R}^k$ to a subset of $\mathbb{R}^d$ by setting $\mathcal{V} = \mathcal{V}_k \times \{0\}^{d-k}$. For a parameter $\delta \in (0, 1/2]$ to be chosen, we define a family of probability distributions $\{P_\nu\}_{\nu \in \mathcal{V}}$ constructively. In particular, the random vector $X \sim P_\nu$ (a single observation) is formed by the following procedure:

$$\text{Choose index } j \in \{1, \ldots, k\} \text{ uniformly at random and set } X = \begin{cases} re_j & \text{w.p. } \frac{1+\delta\nu_j}{2} \\ -re_j & \text{w.p. } \frac{1-\delta\nu_j}{2}. \end{cases} \tag{51}$$

By construction, these distributions have mean vectors

$$\theta_\nu := \mathbb{E}_{P_\nu}[X] = \frac{\delta r}{k}\nu.$$

Consequently, given the properties of the packing $\mathcal{V}$, we have $X \in \mathbb{B}_1(r)$ with probability 1, and fixing $t \in \mathbb{R}_+$, we have that the associated separation function (20) satisfies

$$\delta^2(t) \geq \min\left\{\|\theta_\nu - \theta_{\nu'}\|_2^2 \mid \left\|\nu - \nu'\right\|_1 \geq t\right\} \geq \frac{r^2\delta^2}{k^2} \min\left\{\left\|\nu - \nu'\right\|_2^2 \mid \left\|\nu - \nu'\right\|_1 \geq t\right\} \geq \frac{2r^2\delta^2}{k^2}t.$$

We claim that so long as $t \leq k/6$ and $k \geq 3$, we have

$$\log \frac{|\mathcal{V}|}{N_t} > \max\left\{\frac{k}{6}, 2\right\}. \tag{52}$$

Indeed, for $t \in \mathbb{N}$ with $t \leq k/2$, we see by the binomial theorem that

$$N_t = \sum_{\tau=0}^{t} \binom{d}{\tau} \leq 2\binom{k}{t} \leq 2\left(\frac{ke}{t}\right)^t.$$

Consequently, for $t \leq k/6$, the ratio $|\mathcal{V}|/N_t$ satisfies

$$\log \frac{|\mathcal{V}|}{N_t} \geq k \log 2 - \log 2\binom{k}{t} \geq k \log 2 - \frac{k}{6}\log(6e) - \log 2 = k \log \frac{2}{2^{1/k}\sqrt[6]{6e}} > \max\left\{\frac{k}{6}, 2\right\}$$

for $k \geq 12$. The case $2 \leq k < 12$ can be checked directly, yielding claim (52).

Thus we see that the mean vectors $\{\theta_\nu\}_{\nu \in \mathcal{V}}$ provide us with an $r\delta\sqrt{2t}/k$-separated set (in ℓ_2-norm) with log ratio of its size at least $\max\{k/6, 2\}$.

Upper bounding the mutual information: Our next step is to bound the mutual information $I(Z_1, \ldots, Z_n; V)$ when the observations X come from the distribution (51) and V is uniform in the set $\mathcal{V}$. We have the following lemma, which applies so long as the channel Q is non-interactive and α-locally private (4). See Appendix H.1 for the proof.

Lemma 4. *Fix $k \in \{1, \ldots, d\}$. Let Z_i be α-locally differentially private for X_i, and let X be sampled according to the distribution (51) conditional on $V = \nu$. Then*

$$I(Z_1, \ldots, Z_n; V) \le n \frac{\delta^2}{4k} (e^\alpha - 1)^2.$$

Applying testing inequalities: We now show how a combination the sampling scheme (51) and Lemma 4 give us our desired lower bound. Fix $k \le d$ and let $\mathcal{V} = \{-1, 1\}^k \times \{0\}^{d-k}$. Combining Lemma 4 and the fact that the vectors θ_ν provide a $r\delta/\sqrt{2t}/k$-separated set of log-cardinality at least $\max\{k/6, 2\}$, the minimax Fano bound (21) implies that for any $k \in \{1, \ldots, d\}$ and $t \le k/6$, we have

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_2^2, \alpha) \ge \frac{r^2 \delta^2 t}{2k^2} \left(1 - \frac{n\delta^2 (e^\alpha - 1)^2/(4k) + \log 2}{\max\{k/6, 2\}}\right).$$

Because of the one-dimensional mean-estimation lower bounds provided in Section 3.2.1, we may assume w.l.o.g. that $k \ge 12$. Setting $t = k/6$ and $\delta_{n,\alpha,k}^2 = \min\{1, k^2/(3n(e^\alpha - 1)^2)\}$, we obtain

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_2^2, \alpha) \ge \frac{r^2 \delta_{n,\alpha,k}^2}{12k} \left(1 - \frac{1}{2} - \frac{\log 2}{2}\right) \ge \frac{1}{80} r^2 \min\left\{\frac{1}{k}, \frac{k}{3n(e^\alpha - 1)^2}\right\}$$

for a universal (numerical) constant c. Since $(e^\alpha - 1)^2 < 3\alpha^2$ for $\alpha \in [0, 1]$, we obtain the lower bound

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_2^2, \alpha) \ge \frac{1}{80} r^2 \max_{k \in [d]} \left\{\min\left\{\frac{1}{k}, \frac{k}{9n\alpha^2}\right\}\right\}$$

for $\alpha \in [0, 1]$. Setting k in the preceding display to be the integer in $\{1, \ldots, d\}$ nearest $\sqrt{n\alpha^2}$ gives the result of the proposition.

D.2 Proof of Corollary 6

Constructing a well-separated set: In this case, the packing set is very simple: set $\mathcal{V} = \{\pm e_j\}_{j=1}^d$ so that $|\mathcal{V}| = 2d$. Fix some $\delta \in [0, 1]$, and for $\nu \in \mathcal{V}$, define a distribution P_ν supported on $\mathcal{X} = \{-r, r\}^d$ via

$$P_\nu(X = x) = (1 + \delta \nu^\top x / r)/2^d.$$

In words, for $\nu = e_j$, the coordinates of X are independent uniform on $\{-r, r\}$ except for the coordinate j, for which $X_j = r$ with probability $1/2 + \delta\nu_j$ and $X_j = -r$ with probability $1/2 - \delta\nu_j$. With this scheme, we have $\theta(P_\nu) = r\delta\nu$, which is 1-sparse, and since $\|\delta r\nu - \delta r\nu'\|_2 \ge \sqrt{2}\delta r$, we have constructed a $\sqrt{2}\delta r$ packing in ℓ_2-norm.

Upper bounding the mutual information: Let V be drawn uniformly from the packing set $\mathcal{V} = \{\pm e_j\}_{j=1}^d$. With the sampling scheme in the previous paragraph, we may provide the following upper bound on the mutual information $I(Z_1, \ldots, Z_n; V)$ for any non-interactive private distribution (4):

Lemma 5. *For any non-interactive α-differentially private distribution Q, we have*

$$I(Z_1, \ldots, Z_n; V) \le n \frac{e^\alpha}{4d} \left(e^\alpha - e^{-\alpha}\right)^2 \delta^2.$$

See Appendix H.2 for a proof.

Applying testing inequalities: Finally, we turn to application of the testing inequalities. Lemma 5, in conjunction with the standard testing reduction and Fano's inequality (21) with the choice $t = 0$, implies that

$$\mathfrak{M}_n(\theta(\mathcal{P}), \|\cdot\|_2^2, \alpha) \ge \frac{r^2\delta^2}{2} \left(1 - \frac{e^\alpha \delta^2 n (e^\alpha - e^{-\alpha})^2/(4d) + \log 2}{\log(2d)}\right).$$

There is no loss of generality in assuming that $d \ge 2$, in which case the choice

$$\delta^2 = \min\left\{1, \frac{d \log(2d)}{e^\alpha (e^\alpha - e^{-\alpha})^2 n}\right\}$$

yields the lower bound.

It thus remains to provide the upper bound. In this case, we use the sampling strategy (28) of Section 4.2.3, noting that we may take the bound B on $\|Z\|_\infty$ to be $B = c\sqrt{d}r/\alpha$ for a constant c. Let θ^* denote the true mean, assumed to be s-sparse. Now consider estimating θ^* by the ℓ_1-regularized optimization problem

$$\widehat{\theta} := \operatorname*{argmin}_{\theta \in \mathbb{R}^d} \left\{ \frac{1}{2n} \left\| \sum_{i=1}^n (Z_i - \theta) \right\|_2^2 + \lambda \|\theta\|_1 \right\},$$

Defining the error vector $W = \theta^* - \frac{1}{n}\sum_{i=1}^n Z_i$, we claim that

$$\lambda \ge 2\|W\|_\infty \quad \text{implies that} \quad \|\widehat{\theta} - \theta\|_2 \le 3\lambda\sqrt{s}. \tag{53}$$

This result is a consequence of standard results on sparse estimation (e.g., Negahban et al. [39, Theorem 1 and Corollary 1]).

Now we note if $W_i = \theta^* - Z_i$, then $W = \frac{1}{n}\sum_{i=1}^n W_i$, and by construction of the sampling mechanism (28) we have $\|W_i\|_\infty \le c\sqrt{d}r/\alpha$ for a constant c. By Hoeffding's inequality and a union bound, we thus have for some (different) universal constant c that

$$\mathbb{P}(\|W\|_\infty \ge t) \le 2d \exp\left(-c\frac{n\alpha^2 t^2}{r^2 d}\right) \quad \text{for } t \ge 0.$$

By taking $t^2 = r^2 d(\log(2d) + \epsilon^2)/(cn\alpha^2)$, we find that $\|W\|_\infty^2 \le r^2 d(\log(2d) + \epsilon^2)/(cn\alpha^2)$ with probability at least $1 - \exp(-\epsilon^2)$, which gives the claimed minimax upper bound by appropriate choice of $\lambda = c\sqrt{d \log d/n\alpha^2}$ in inequality (53).

D.3 Proof of inequality (26)

We prove the bound by an argument using the private form of Fano's inequality (21) with $t = 0$ and replacing X with Z. The proof makes use of the classical Varshamov-Gilbert bound (e.g. [52, Lemma 4]):

Lemma 6 (Varshamov-Gilbert)**.** *There is a packing $\mathcal{V}$ of the d-dimensional hypercube $\{-1,1\}^d$ of size $|\mathcal{V}| \geq \exp(d/8)$ such that*

$$\left\|\nu - \nu'\right\|_1 \geq d/2 \quad \text{for all distinct pairs } \nu, \nu' \in \mathcal{V}.$$

Now, let $\delta \in [0,1]$ and the distribution P_ν be a point mass at $\delta\nu/\sqrt{d}$. Then $\theta(P_\nu) = \delta\nu/\sqrt{d}$ and $\|\theta(P_\nu) - \theta(P_{\nu'})\|_2^2 \geq \delta^2$. In addition, a calculation implies that if M_1 and M_2 are d-dimensional Laplace(κ) distributions with means θ_1 and θ_2, respectively, then

$$D_{\rm kl}\left(M_1 \| M_2\right) = \sum_{j=1}^d \left(\exp(-\kappa|\theta_{1,j} - \theta_{2,j}|) + \kappa|\theta_{1,j} - \theta_{2,j}| - 1\right) \leq \frac{\kappa^2}{2}\left\|\theta_1 - \theta_2\right\|_2^2.$$

As a consequence, we have that under our Laplacian sampling scheme for the Z and with V chosen uniformly from $\mathcal{V}$,

$$I(Z_1, \ldots, Z_n; V) \leq \frac{1}{|\mathcal{V}|^2} n \sum_{\nu,\nu' \in \mathcal{V}} D_{\rm kl}\left(M_\nu \| M_{\nu'}\right) \leq \frac{n\alpha^2}{2d|\mathcal{V}|^2} \sum_{\nu,\nu' \in \mathcal{V}} \left\|(\delta/\sqrt{d})(\nu - \nu')\right\|_2^2 \leq \frac{2n\alpha^2\delta^2}{d}.$$

Now, applying Fano's inequality (21), we find that

$$\inf_{\widehat{\theta}} \sup_{\nu \in \mathcal{V}} \mathbb{E}_{P_\nu}\left[\|\widehat{\theta}(Z_1, \ldots, Z_n) - \theta(P_\nu)\|_2^2\right] \geq \frac{\delta^2}{4}\left(1 - \frac{2n\alpha^2\delta^2/d + \log 2}{d/8}\right).$$

We may assume (based on our one-dimensional results in Corollary 1) w.l.o.g. that $d \geq 10$. Taking $\delta^2 = d^2/(48n\alpha^2)$ then implies the result (26).

E Proof of Theorem 3

The proof of this theorem combines the techniques we used in the proofs of Theorems 1 and 2; the first handles interactivity, while the techniques to derive the variational bounds are reminiscent of those used in Theorem 2. Our first step is to note a consequence of the independence structure in Fig. 2 essential to our tensorization steps. More precisely, we claim that for any set $S \in \sigma(\mathcal{Z})$,

$$M_{\pm j}(Z_i \in S \mid z_{1:i-1}) = \int Q(Z_i \in S \mid Z_{1:i-1} = z_{1:i-1}, X_i = x) dP_{\pm j,i}(x). \tag{54}$$

We postpone the proof of this intermediate claim to the end of this section.

Now consider the summed KL-divergences. Let $M_{\pm j,i}(\cdot \mid z_{1:i-1})$ denote the conditional distribution of Z_i under $P_{\pm j}$, conditional on $Z_{1:i-1} = z_{1:i-1}$. As in the proof of Corollary 4, the chain rule for KL-divergences [e.g., 26, Chapter 5] implies

$$D_{\rm kl}\left(M_{+j}^n \| M_{-j}^n\right) = \sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} D_{\rm kl}\left(M_{+j}(\cdot \mid z_{1:i-1}) \| M_{-j}(\cdot \mid z_{1:i-1})\right) dM_{+j}^{i-1}(z_{1:i-1}).$$

For notational convenience in the remainder of the proof, let us recall that the symmetrized KL divergence between measures M and M' is $D_{\rm kl}^{\rm sy}(M\|M') = D_{\rm kl}(M\|M') + D_{\rm kl}(M'\|M)$.

Defining $\overline{P} := 2^{-d}\sum_{\nu\in\mathcal{V}} P_\nu^n$, we have $2\overline{P} = P_{+j} + P_{-j}$ for each j simultaneously, We also introduce $\overline{M}(S) = \int Q(S \mid x_{1:n}) d\overline{M}(x_{1:n})$, and let $\mathbb{E}_{\pm j}$ denote the expectation taken under the marginals $M_{\pm j}$. We then have

$$\begin{aligned}
&D_{\rm kl}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\rm kl}\left(M_{-j}^n \| M_{+j}^n\right) \\
&= \sum_{i=1}^n \left(\mathbb{E}_{+j}[D_{\rm kl}\left(M_{+j,i}(\cdot \mid Z_{1:i-1}) \| M_{-j,i}(\cdot \mid Z_{1:i-1})\right)] + \mathbb{E}_{-j}[D_{\rm kl}\left(M_{-j,i}(\cdot \mid Z_{1:i-1}) \| M_{+j,i}(\cdot \mid Z_{1:i-1})\right)]\right) \\
&\le \sum_{i=1}^n \left(\mathbb{E}_{+j}[D_{\rm kl}^{\rm sy}\left(M_{+j,i}(\cdot \mid Z_{1:i-1}) \| M_{-j,i}(\cdot \mid Z_{1:i-1})\right)] + \mathbb{E}_{-j}[D_{\rm kl}^{\rm sy}\left(M_{+j,i}(\cdot \mid Z_{1:i-1}) \| M_{-j,i}(\cdot \mid Z_{1:i-1})\right)]\right) \\
&= 2\sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} D_{\rm kl}^{\rm sy}\left(M_{+j,i}(\cdot \mid z_{1:i-1}) \| M_{-j,i}(\cdot \mid z_{1:i-1})\right) d\overline{M}^{i-1}(z_{1:i-1}),
\end{aligned}$$

where we have used the definition of $\overline{M}$ and that $2\overline{P} = P_{+j} + P_{-j}$ for all j. Summing over $j \in [d]$ yields

$$\sum_{j=1}^d D_{\rm kl}^{\rm sy}\left(M_{+j}^n \| M_{-j}^n\right) \le 2\sum_{i=1}^n \int_{\mathcal{Z}^{i-1}} \sum_{j=1}^d \underbrace{D_{\rm kl}^{\rm sy}\left(M_{+j,i}(\cdot \mid z_{1:i-1}) \| M_{-j,i}(\cdot \mid z_{1:i-1})\right)}_{=:\mathcal{T}_{j,i}} d\overline{M}^{i-1}(z_{1:i-1}). \quad (55)$$

We bound the underlined expression in inequality (55), whose elements we denote by $\mathcal{T}_{j,i}$.

Without loss of generality (as in the proof of Theorem 2), we may assume $\mathcal{Z}$ is finite, and that $\mathcal{Z} = \{1, 2, \ldots, k\}$ for some positive integer k. Using the probability mass functions $m_{\pm j,i}$ and omitting the index i when it is clear from context, Lemma 3 implies

$$\begin{aligned}
\mathcal{T}_{j,i} &= \sum_{z=1}^k \left(m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1})\right) \log \frac{m_{+j}(z \mid z_{1:i-1})}{m_{-j}(z \mid z_{1:i-1})} \\
&\le \sum_{z=1}^k \left(m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1})\right)^2 \frac{1}{\min\{m_{+j}(z \mid z_{1:i-1}), m_{-j}(z \mid z_{1:i-1})\}}.
\end{aligned}$$

For each fixed $z_{1:i-1}$, define the infimal measure $m^0(z \mid z_{1:i-1}) := \inf_{x\in\mathcal{X}} q(z \mid X_i = x, z_{1:i-1})$. By construction, we have $\min\{m_{+j}(z \mid z_{1:i-1}), m_{-j}(z \mid z_{1:i-1})\} \ge m^0(z \mid z_{1:i-1})$, and hence

$$\mathcal{T}_{j,i} \le \sum_{z=1}^k \left(m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1})\right)^2 \frac{1}{m^0(z \mid z_{1:i-1})}.$$

Recalling equality (54), we have

$$\begin{aligned}
m_{+j}(z \mid z_{1:i-1}) - m_{+j}(z \mid z_{1:i-1}) &= \int_{\mathcal{X}} q(z \mid x, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x)) \\
&= m^0(z \mid z_{1:i-1}) \int_{\mathcal{X}} \left(\frac{q(z \mid x, z_{1:i-1})}{m^0(z \mid z_{1:i-1})} - 1\right)(dP_{+j,i}(x) - dP_{-j,i}(x)).
\end{aligned}$$

From this point, the proof is similar to that of Theorem 2. Define the collection of functions

$$\mathcal{F}_\alpha := \{f : \mathcal{X} \times \mathcal{Z}^i \to [0, e^\alpha - 1]\}.$$

Using the definition of differential privacy, we have $\frac{q(z|x,z_{1:i-1})}{m^0(z|z_{1:i-1})} \in [1, e^\alpha]$, so there exists $f \in \mathcal{F}_\alpha$ such that

$$\begin{aligned}\sum_{j=1}^d \mathcal{T}_{j,i} &\le \sum_{j=1}^d \sum_{z=1}^k \frac{\left(m^0(z \mid z_{1:i-1})\right)^2}{m^0(z \mid z_{1:i-1})} \left(\int_{\mathcal{X}} f(x, z, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x)) \right)^2 \\ &= \sum_{z=1}^k m^0(z \mid z_{1:i-1}) \sum_{j=1}^d \left(\int_{\mathcal{X}} f(x, z, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x)) \right)^2.\end{aligned}$$

Taking a supremum over $\mathcal{F}_\alpha$, we find the further upper bound

$$\sum_{j=1}^d \mathcal{T}_{j,i} \le \sum_{z=1}^k m^0(z \mid z_{1:i-1}) \sup_{f \in \mathcal{F}_\alpha} \sum_{j=1}^d \left(\int_{\mathcal{X}} f(x, z, z_{1:i-1})(dP_{+j,i}(x) - dP_{-j,i}(x)) \right)^2.$$

The inner supremum may be taken independently of z and $z_{1:i-1}$, so we rescale by $(e^\alpha - 1)$ to obtain our penultimate inequality

$$\begin{aligned}&\sum_{j=1}^d D_{\rm kl}^{\rm sy}\left(M_{+j,i}(\cdot \mid z_{1:i-1}) \| M_{-j,i}(\cdot \mid z_{1:i-1})\right) \\ &\qquad \le (e^\alpha - 1)^2 \sum_{z=1}^k m^0(z \mid z_{1:i-1}) \sup_{\gamma \in \mathbb{B}_\infty(\mathcal{X})} \sum_{j=1}^d \left(\int_{\mathcal{X}} \gamma(x)(dP_{+j,i}(x) - dP_{-j,i}(x)) \right)^2.\end{aligned}$$

Noting that m^0 sums to a quantity ≤ 1 and substituting the preceding expression in inequality (55) completes the proof.

Finally, we return to prove our intermediate marginalization claim (54). We have that

$$\begin{aligned}M_{\pm j}(Z_i \in S \mid z_{1:i-1}) &= \int Q(Z_i \in S \mid z_{1:i-1}, x_{1:n}) dP_{\pm j}(x_{1:n} \mid z_{1:i-1}) \\ &\stackrel{(i)}{=} \int Q(Z_i \in S \mid z_{1:i-1}, x_i) dP_{\pm j}(x_{1:n} \mid z_{1:i-1}) \\ &\stackrel{(ii)}{=} \int Q(Z_i \in S \mid Z_{1:i-1} = z_{1:i-1}, X_i = x) dP_{\pm j,i}(x),\end{aligned}$$

where equality (i) follows by the assumed conditional independence structure of Q (recall Figure 2) and equality (ii) is a consequence of the independence of X_i and $Z_{1:i-1}$ under $P_{\pm j}$. That is, we have $P_{+j}(X_i \in S \mid Z_{1:i-1} = z_{1:i-1}) = P_{+j,i}(S)$ by the definition of P_ν^n as a product and that $P_{\pm j}$ are a mixture of the products P_ν^n.

F Proofs of multinomial estimation results

In this section, we prove the lower bounds in Corollary 7. Before proving the bounds, however, we outline our technique, which borrows from that in Section D, and which we also use to prove the lower bounds on density estimation. The outline is as follows:

(1) As in step (1) of Section D, our first step is a standard reduction using the sharper version of Assouad's method (Lemma 1) from estimation to a multiple binary hypothesis testing problem. Specifically, we perform a (essentially standard) reduction of the form (29).

(2) Having constructed appropriately separated binary hypothesis tests, we use apply Theorem 3 via inequality (41) to control the testing error in the binary testing problem. Applying the theorem requires bounding certain suprema related to the covariance structure of randomly selected elements of $\mathcal{V} = \{-1, 1\}^d$, as in the arguments in Section D. This is made easier by the symmetry of the binary hypothesis testing problems.

With this outline in mind, we turn to the proofs of inequalities (32a) and (32b). As we proved the upper bounds in Section 5.2.1, this section focuses on the argument for the lower bound. We provide the full proof for the mean-squared Euclidean error, after which we show how the result for the ℓ_1-error follows.

Our first step is to provide a lower bound of the form (29), giving a Hamming separation for the squared error. To that end, fix $\delta \in [0, 1]$, and for simplicity, let us assume that d is even. In this case, we set $\mathcal{V} = \{-1, 1\}^{d/2}$, and for $\nu \in \mathcal{V}$ let P_ν be the multinomial distribution with parameter

$$\theta_\nu := \frac{1}{d}\mathbb{1} + \delta\frac{1}{d}\begin{bmatrix}\nu \\ -\nu\end{bmatrix} \in \Delta_d.$$

For any estimator $\widehat{\theta}$, by defining $\widehat{\nu}_j = \operatorname{sign}(\widehat{\theta}_j - 1/d)$ for $j \in [d/2]$ we have the lower bound

$$\|\widehat{\theta} - \theta_\nu\|_2^2 \geq \frac{\delta^2}{d^2}\sum_{j=1}^{d/2}\mathbf{1}\{\widehat{\nu}_j \neq \nu_j\},$$

so that by the sharper variant (41) of Assouad's Lemma, we obtain

$$\max_{\nu\in\mathcal{V}}\mathbb{E}_{P_\nu}[\|\widehat{\theta} - \theta_\nu\|_2^2] \geq \frac{\delta^2}{4d}\left[1 - \left(\frac{1}{2d}\sum_{j=1}^{d/2} D_{\rm kl}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\rm kl}\left(M_{-j}^n \| M_{+j}^n\right)\right)^{\frac{1}{2}}\right]. \tag{56}$$

Now we apply Theorem 3, which requires bounding sums of integrals $\int \gamma(dP_{+j} - dP_{-j})$, where P_{+j} is defined in expression (30). We claim the following inequality:

$$\sup_{\gamma\in\mathbb{B}_\infty(\mathcal{X})}\sum_{j=1}^{d/2}\left(\int_{\mathcal{X}}\gamma(x)dP_{+j}(x) - dP_{-j}(x)\right)^2 \leq \frac{8\delta^2}{d}. \tag{57}$$

Indeed, by construction P_{+j} is the multinomial with parameter $(1/d)\mathbb{1} + (\delta/d)[e_j^\top \quad -e_j^\top]^\top \in \Delta_d$ and similarly for P_{-j}, where $e_j \in \{0, 1\}^{d/2}$ denotes the jth standard basis vector. Abusing notation and identifying γ with vectors $\gamma \in [-1, 1]^d$, we have

$$\int_{\mathcal{X}}\gamma(x)dP_{+j}(x) - dP_{-j}(x) = \frac{2\delta}{d}\gamma^\top\begin{bmatrix}e_j \\ -e_j\end{bmatrix},$$

whence we find

$$\sum_{j=1}^{d/2}\left(\int_{\mathcal{X}}\gamma(x)dP_{+j}(x)-dP_{-j}(x)\right)^2=\frac{4\delta^2}{d^2}\gamma^\top\sum_{j=1}^{d/2}\begin{bmatrix}e_j\\-e_j\end{bmatrix}\begin{bmatrix}e_j\\-e_j\end{bmatrix}^\top\gamma=\frac{4\delta^2}{d^2}\gamma^\top\begin{bmatrix}I&-I\\-I&I\end{bmatrix}\gamma\le\frac{8\delta^2}{d},$$

because the operator norm of the matrix is bounded by 2. This gives the claim (57).

Substituting the bound (57) into the bound (56) via Theorem 3, we obtain

$$\max_{\nu\in\mathcal{V}}\mathbb{E}_{P_\nu}[\|\widehat{\theta}-\theta_\nu\|_2^2]\ge\frac{\delta^2}{4d}\left[1-\left(4n(e^\alpha-1)^2\delta^2/d^2\right)^{\frac{1}{2}}\right].$$

Choosing $\delta^2=\min\{1,d^2/(16n(e^\alpha-1)^2)\}$ gives the lower bound

$$\mathfrak{M}_n(\Delta_d,\|\cdot\|_2^2,\alpha)\ge\min\left\{\frac{1}{4d},\frac{d}{64n(e^\alpha-1)^2}\right\}.$$

To complete the proof, we note that we can prove the preceding upper bound for any even $d_0\in\{2,\ldots,d\}$; this requires choosing $\nu\in\mathcal{V}=\{-1,1\}^{d_0/2}$ and constructing the multinomial vectors

$$\theta_\nu=\frac{1}{d_0}\begin{bmatrix}\mathbb{1}_{d_0}\\0_{d-d_0}\end{bmatrix}+\frac{\delta}{d_0}\begin{bmatrix}\nu\\-\nu\\0_{d-d_0}\end{bmatrix}\in\Delta_d,\quad\text{where }\mathbb{1}_{d_0}=[1\;1\;\cdots\;1]^\top\in\mathbb{R}^{d_0}.$$

Repeating the proof *mutatis mutandis* gives the bound

$$\mathfrak{M}_n(\Delta_d,\|\cdot\|_2^2,\alpha)\ge\max_{d_0\in\{2,4,\ldots,2\lfloor d/2\rfloor\}}\min\left\{\frac{1}{4d_0},\frac{d_0}{64n(e^\alpha-1)^2}\right\}.$$

Choosing d_0 to be the even integer closest to $\sqrt{n\alpha^2}$ in $\{1,\ldots,d\}$ and noting that $(e^\alpha-1)^2\le3\alpha^2$ for $\alpha\in[0,1]$ gives the claimed result (32a).

In the case of measuring error in the ℓ_1-norm, we provide a completely identical proof, except that we have the separation $\|\widehat{\theta}-\theta_\nu\|_1\ge(\delta/d)\sum_{j=1}^{d/2}\mathbf{1}\{\widehat{\nu}_j\ne\nu_j\}$, and thus inequality (56) holds with the initial multiplier $\delta^2/(4d)$ replaced by $\delta/(4d)$. Parallel reasoning to the ℓ_2^2 case then gives the minimax lower bound

$$\mathfrak{M}_n(\Delta_d,\|\cdot\|_1,\alpha)\ge\frac{\delta}{4d_0}\left[1-\left(4n(e^\alpha-1)^2\delta^2/d_0^2\right)^{\frac{1}{2}}\right]$$

for any even $d_0\in\{2,\ldots,d\}$. Choosing $\delta=\min\{1,d_0^2/(16n(e^\alpha-1)^2)\}$ gives the claim (32b).

G Proof of Corollary 8

In this section, we provide the proof of Corollary 8 on density esitmation from Section 5.2.2. We defer the proofs of more technical results to later appendices. We provide the proof of the minimax lower bound in Section G.1 and the upper bound in Section G.2. Throughout all proofs, we use c to denote a constant whose value may change from line to line.

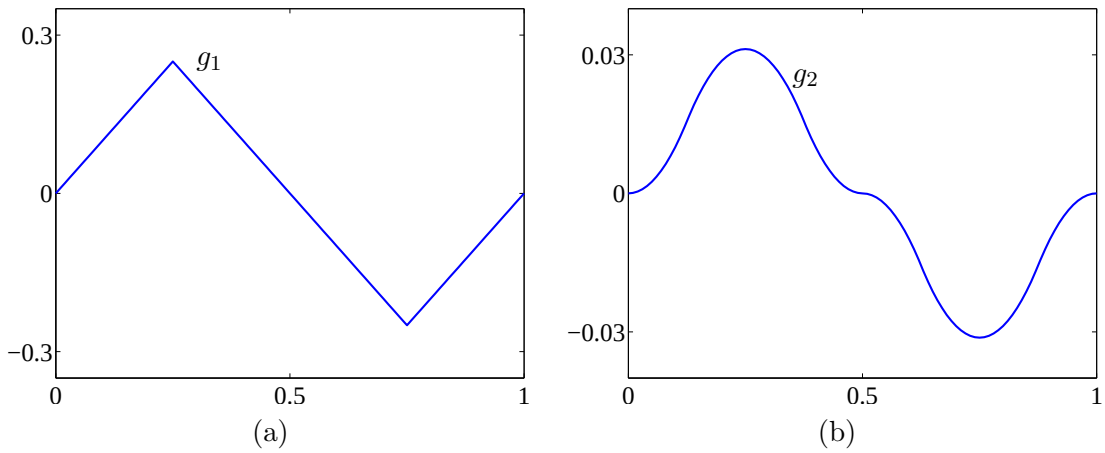

Figure 9. Panel (a): illustration of 1-Lipschitz continuous bump function g_1 used to pack $\mathcal{F}_\beta$ when $\beta = 1$. Panel (b): bump function g_2 with $|g_2''(x)| \leq 1$ used to pack $\mathcal{F}_\beta$ when $\beta = 2$.

G.1 Proof of the lower bound (37)

As with our proof for multinomial estimation, the argument follows the general outline described at the beginning of Section F. We remark that our proof is based on an explicit construction of densities identified with corners of the hypercube, a more classical approach than the global metric entropy approach of Yang and Barron [51] (cf. [52]). We use the local packing approach since it is better suited to the privacy constraints and information contractions that we have developed. In comparison with our proofs of previous propositions, the construction of a suitable packing of $\mathcal{F}_\beta$ is somewhat more challenging: the identification of densities with finite-dimensional vectors, which we require for our application of Theorem 3, is not immediately obvious. In all cases, we guarantee that our density functions f belong to the trigonometric Sobolev space, so we may work directly with smooth density functions f.

Constructing well-separated densities: We begin by describing a standard framework for defining local packings of density functions. Let $g_\beta : [0,1] \to \mathbb{R}$ be a function satisfying the following properties:

(a) The function g_β is β-times differentiable with

$$0 = g_\beta^{(i)}(0) = g_\beta^{(i)}(1/2) = g_\beta^{(i)}(1) \quad \text{for all } i < \beta.$$

(b) The function g_β is centered with $\int_0^1 g_\beta(x)dx = 0$, and there exist constants $c, c_{1/2} > 0$ such that

$$\int_0^{1/2} g_\beta(x)dx = -\int_{1/2}^1 g_\beta(x)dx = c_{1/2} \quad \text{and} \quad \int_0^1 \left(g_\beta^{(i)}(x)\right)^2 dx \geq c \quad \text{for all } i < \beta.$$

(c) The function g_β is nonnegative on $[0, 1/2]$ and non-positive on $[1/2, 1]$, and Lebesgue measure is absolutely continuous with respect to the measures $G_j, j = 1, 2$, given by

$$G_1(A) = \int_{A\cap[0,1/2]} g_\beta(x)dx \quad \text{and} \quad G_2(A) = -\int_{A\cap[1/2,1]} g_\beta(x)dx. \tag{58}$$

(d) Lastly, for almost every $x \in [0,1]$, we have $|g_\beta^{(\beta)}(x)| \le 1$ and $|g_\beta(x)| \le 1$.

As illustrated in Figure 9, the functions g_β are smooth "bump" functions.

Fix a positive integer k (to be specified in the sequel). Our first step is to construct a family of "well-separated" densities for which we can reduce the density estimation problem to one of identifying corners of a hypercube, which allows application of Lemma 1. Specifically, we must exhibit a condition similar to the separation condition (29). For each $j \in \{1, \ldots, k$ define the function

$$g_{\beta,j}(x) := \frac{1}{k^\beta} g_\beta\left(k\Big(x - \frac{j-1}{k}\Big)\right) \mathbf{1}\Big\{x \in \Big[\tfrac{j-1}{k}, \tfrac{j}{k}\Big]\Big\}.$$

Based on this definition, we define the family of densities

$$\left\{f_\nu := 1 + \sum_{j=1}^k \nu_j g_{\beta,j} \quad \text{for } \nu \in \mathcal{V}\right\} \subseteq \mathcal{F}_\beta. \tag{59}$$

It is a standard fact [52, 47] that for any $\nu \in \mathcal{V}$, the function f_ν is β-times differentiable, satisfies $|f^{(\beta)}(x)| \le 1$ for all x. Now, based on some density $f \in \mathcal{F}_\beta$, let us define the sign vector $\mathsf{v}(f) \in \{-1,1\}^k$ to have entries

$$\mathsf{v}_j(f) := \operatorname*{argmin}_{s \in \{-1,1\}} \int_{[\frac{j-1}{k}, \frac{j}{k}]} (f(x) - s g_{\beta,j}(x))^2 \, dx.$$

Then by construction of the g_β and v, we have for a numerical constant c (whose value may depend on β) that

$$\|f - f_\nu\|_2^2 \ge c \sum_{j=1}^k \mathbf{1}\{\mathsf{v}_j(f) \ne \nu_j\} \int_{[\frac{j-1}{k}, \frac{j}{k}]} (g_{\beta,j}(x))^2 dx = \frac{c}{k^{2\beta+1}} \sum_{j=1}^k \mathbf{1}\{\mathsf{v}_j(f) \ne \nu_j\}.$$

By inspection, this is the Hamming separation required in inequality (29), whence the sharper version (41) of Assouad's Lemma 1 gives the result

$$\mathfrak{M}_n\left(\mathcal{F}_\beta[1], \|\cdot\|_2^2, \alpha\right) \ge \frac{c}{k^{2\beta}} \left[1 - \left(\frac{1}{4k} \sum_{j=1}^k (D_{\rm kl}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\rm kl}\left(M_{-j}^n \| M_{+j}^n\right))\right)^{\frac{1}{2}}\right], \tag{60}$$

where we have defined $P_{\pm j}$ to be the probability distribution associated with the averaged densities $f_{\pm j} = 2^{1-k} \sum_{\nu : \nu_j = \pm 1} f_\nu$.

Applying divergence inequalities: Now we must control the summed KL-divergences. To do so, we note that by the construction (59), symmetry implies that

$$f_{+j} = 1 + g_{\beta,j} \quad \text{and} \quad f_{-j} = 1 - g_{\beta,j} \quad \text{for each} \quad j \in [k]. \tag{61}$$

We then obtain the following result, which bounds the averaged KL-divergences.

Lemma 7. *For any α-locally private conditional distribution Q, the summed KL-divergences are bounded as*

$$\sum_{j=1}^{k} \left(D_{\mathrm{kl}} \left(M_{+j}^n \| M_{-j}^n \right) + D_{\mathrm{kl}} \left(M_{+j}^n \| M_{-j}^n \right) \right) \leq 4 c_{1/2}^2 \, n \frac{(e^\alpha - 1)^2}{k^{2\beta+1}}.$$

The proof of this lemma is fairly involved, so we defer it to Appendix H.3. We note that, for $\alpha \leq 1$, we have $(e^\alpha - 1)^2 \leq 3\alpha^2$, so we may replace the bound in Lemma 7 with the quantity $cn\alpha^2/k^{2\beta+1}$ for a constant c. We remark that standard divergence bounds using Assouad's lemma [52, 47] provide a bound of roughly $n/k^{2\beta}$; our bound is thus essentially a factor of the "dimension" k tighter.

The remainder of the proof is an application of inequality (60). In particular, by applying Lemma 7, we find that for any α-locally private channel Q, there are constants c_0, c_1 (whose values may depend on β) such that

$$\mathfrak{M}_n \left(\mathcal{F}_\beta, \|\cdot\|_2^2, Q \right) \geq \frac{c_0}{k^{2\beta}} \left[1 - \left(\frac{c_1 n \alpha^2}{k^{2\beta+2}} \right)^{\frac{1}{2}} \right].$$

Choosing $k_{n,\alpha,\beta} = \left(4 c_1 n \alpha^2\right)^{\frac{1}{2\beta+2}}$ ensures that the quantity inside the parentheses is at least $1/2$. Substituting for k in the preceding display proves the proposition.

G.2 Proof of the upper bound (37)

We begin by fixing $k \in \mathbb{N}$; we will optimize the choice of k shortly. Recall that, since $f \in \mathcal{F}_\beta[r]$, we have $f = \sum_{j=1}^\infty \theta_j \varphi_j$ for $\theta_j = \int f \varphi_j$. Thus we may define $\overline{Z}_j = \frac{1}{n} \sum_{i=1}^n Z_{i,j}$ for each $j \in \{1, \ldots, k\}$, and we have

$$\|\widehat{f} - f\|_2^2 = \sum_{j=1}^{k} (\theta_j - \overline{Z}_j)^2 + \sum_{j=k+1}^{\infty} \theta_j^2.$$

Since $f \in \mathcal{F}_\beta[r]$, we are guaranteed that $\sum_{j=1}^\infty j^{2\beta} \theta_j^2 \leq r^2$, and hence

$$\sum_{j>k} \theta_j^2 = \sum_{j>k} j^{2\beta} \frac{\theta_j^2}{j^{2\beta}} \leq \frac{1}{k^{2\beta}} \sum_{j>k} j^{2\beta} \theta_j^2 \leq \frac{1}{k^{2\beta}} r^2.$$

For the indices $j \leq k$, we note that by assumption, $\mathbb{E}[Z_{i,j}] = \int \varphi_j f = \theta_j$, and since $|Z_{i,j}| \leq B$, we have

$$\mathbb{E}\left[(\theta_j - \overline{Z}_j)^2 \right] = \frac{1}{n} \operatorname{Var}(Z_{1,j}) \leq \frac{B^2}{n} = \frac{B_0^2}{c_k} \frac{k}{n} \left(\frac{e^\alpha + 1}{e^\alpha - 1} \right)^2,$$

where $c_k = \Omega(1)$ is the constant in expression (39). Putting together the pieces, the mean-squared L^2-error is upper bounded as

$$\mathbb{E}_f \left[\|\widehat{f} - f\|_2^2 \right] \leq c \left(\frac{k^2}{n\alpha^2} + \frac{1}{k^{2\beta}} \right),$$

where c is a constant depending on B_0, c_k, and r. Choose $k = (n\alpha^2)^{1/(2\beta+2)}$ to complete the proof.

H Information bounds

In this appendix, we collect the proofs of lemmas providing mutual information and KL-divergence bounds.

H.1 Proof of Lemma 4

Our strategy is to apply Theorem 2 to bound the mutual information. Without loss of generality, we may assume that $r = 1$ so the set $\mathcal{X} = \{\pm e_j\}_{j=1}^k$, where $e_j \in \mathbb{R}^d$. Thus, under the notation of Theorem 2, we may identify vectors $\gamma \in L^\infty(\mathcal{X})$ by vectors $\gamma \in \mathbb{R}^{2k}$. Noting that $\overline{\nu} = \frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}} \nu = 0$ is the mean element of the "packing" set by our construction, the linear functional φ_ν defined in Theorem 2 is

$$\varphi_\nu(\gamma) = \frac{1}{2k}\sum_{j=1}^k \left[\frac{\delta}{2}\gamma(e_j)\nu_j - \frac{\delta}{2}\gamma(-e_j)\nu_j\right] = \frac{\delta}{4k}\gamma^\top \begin{bmatrix} I_{k\times k} & 0_{k\times d-k} \\ -I_{k\times k} & 0_{k\times d-k}\end{bmatrix}\nu.$$

Define the matrix

$$A := \begin{bmatrix} I_{k\times k} & 0_{k\times d-k} \\ -I_{k\times k} & 0_{k\times d-k}\end{bmatrix} \in \{-1,0,1\}^{2k\times d}.$$

Then we have that

$$\begin{aligned}\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\varphi_\nu(\gamma)^2 &= \frac{\delta^2}{(4k)^2}\gamma^\top A\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\nu\nu^\top A^\top\gamma = \frac{\delta^2}{(4k)^2}\gamma^\top A\operatorname{Cov}(V)A^\top\gamma \\ &= \frac{\delta^2}{(4k)^2}\gamma^\top AA^\top\gamma = \left(\frac{\delta}{4k}\right)^2\gamma^\top\begin{bmatrix} I_{k\times k} & -I_{k\times k} \\ -I_{k\times k} & I_{k\times k}\end{bmatrix}\gamma. \end{aligned} \tag{62}$$

Here we have used that $A\operatorname{Cov}(V)A^\top = AI_{d\times d}A^\top$ by the fact that $\mathcal{V} = \{-1,1\}^k \times \{0\}^{d-k}$.

We complete our proof using the bound (62). The operator norm of the matrix specified in (62) is 2. As a consequence, since we have the containment

$$\mathbb{B}_\infty = \left\{\gamma\in\mathbb{R}^{2k} : \|\gamma\|_\infty \le 1\right\} \subset \left\{\gamma\in\mathbb{R}^{2k} : \|\gamma\|_2^2 \le 2k\right\},$$

we have the inequality

$$\sup_{\gamma\in\mathbb{B}_\infty}\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}}\varphi_\nu(\gamma)^2 \le \frac{\delta^2}{16k^2}\cdot 2\cdot 2k = \frac{1}{4}\frac{\delta^2}{k}.$$

Applying Theorem 2 completes the proof.

H.2 Proof of Lemma 5

It is no loss of generality to assume the radius $r = 1$. We use the notation of Theorem 2, recalling the linear functionals $\varphi_\nu : L^\infty(\mathcal{X}) \to \mathbb{R}$. Because the set $\mathcal{X} = \{-1,1\}^d$, we can identify vectors $\gamma \in L^\infty(\mathcal{X})$ with vectors $\gamma \in \mathbb{R}^{2^d}$. Moreover, we have (by construction) that

$$\begin{aligned}\varphi_\nu(\gamma) &= \sum_{x\in\{-1,1\}^d}\gamma(x)p_\nu(x) - \sum_{x\in\{-1,1\}^d}\gamma(x)\overline{p}(x) \\ &= \frac{1}{2^d}\sum_{x\in\mathcal{X}}\gamma(x)(1+\delta\nu^\top x - 1) = \frac{\delta}{2^d}\sum_{x\in\mathcal{X}}\gamma(x)\nu^\top x.\end{aligned}$$

For each $\nu \in \mathcal{V}$, we may construct a vector $u_\nu \in \{-1,1\}^{2^d}$, indexed by $x \in \{-1,1\}^d$, with

$$u_\nu(x) = \nu^\top x = \begin{cases} 1 & \text{if } \nu = \pm e_j \text{ and } \operatorname{sign}(\nu_j) = \operatorname{sign}(x_j) \\ -1 & \text{if } \nu = \pm e_j \text{ and } \operatorname{sign}(\nu_j) \neq \operatorname{sign}(x_j). \end{cases}$$

For $\nu = e_j$, we see that $u_{e_1}, \ldots, u_{e_d}$ are the first d columns of the standard Hadamard transform matrix (and u_{-e_j} are their negatives). Then we have that $\sum_{x \in \mathcal{X}} \gamma(x) \nu^\top x = \gamma^\top u_\nu$, and

$$\varphi_\nu(\gamma) = \gamma^\top u_\nu u_\nu^\top \gamma.$$

Note also that $u_\nu u_\nu^\top = u_{-\nu} u_{-\nu}^\top$, and as a consequence we have

$$\sum_{\nu \in \mathcal{V}} \varphi_\nu(\gamma)^2 = \frac{\delta^2}{4^d} \gamma^\top \sum_{\nu \in \mathcal{V}} u_\nu u_\nu^\top \gamma = \frac{2\delta^2}{4^d} \gamma^\top \sum_{j=1}^d u_{e_j} u_{e_j}^\top \gamma. \tag{63}$$

But now, studying the quadratic form (63), we note that the vectors u_{e_j} are orthogonal. As a consequence, the vectors (up to scaling) u_{e_j} are the only eigenvectors corresponding to positive eigenvalues of the positive semidefinite matrix $\sum_{j=1}^d u_{e_j} u_{e_j}^\top$. Thus, since the set

$$\mathbb{B}_\infty = \left\{ \gamma \in \mathbb{R}^{2^d} : \|\gamma\|_\infty \leq 1 \right\} \subset \left\{ \gamma \in \mathbb{R}^{2^d} : \|\gamma\|_2^2 \leq 2^d \right\},$$

we have via an eigenvalue calculation that

$$\begin{aligned} \sup_{\gamma \in \mathbb{B}_\infty} \sum_{\nu \in \mathcal{V}} \varphi_\nu(\gamma)^2 &\leq \frac{2\delta^2}{4^d} \sup_{\gamma : \|\gamma\|_2^2 \leq 2^d} \gamma^\top \sum_{j=1}^d u_{e_j} u_{e_j}^\top \gamma \\ &= \frac{2\delta^2}{4^d} \|u_{e_1}\|_2^4 = 2\delta^2, \end{aligned}$$

since $\|u_{e_j}\|_2^2 = 2^d$ for each j. Applying Theorem 2 completes the proof.

H.3 Proof of Lemma 7

This result relies on Theorem 3, along with a careful argument to understand the extreme points of $\gamma \in L^\infty([0,1])$ that we use when applying the result. First, we take the packing $\mathcal{V} = \{-1,1\}^\beta$ and densities f_ν for $\nu \in \mathcal{V}$ as in the construction (59). Overall, our first step is to show for the purposes of applying Theorem 3, it is no loss of generality to identify $\gamma \in L^\infty([0,1])$ with vectors $\gamma \in \mathbb{R}^{2k}$, where γ is constant on intervals of the form $[i/2k, (i+1)/2k]$. With this identification complete, we can then provide a bound on the correlation of any $\gamma \in \mathbb{B}_\infty$ with the densities $f_{\pm j}$ defined in (61), which completes the proof.

With this outline in mind, let the sets D_i, $i \in \{1, 2, \ldots, 2k\}$, be defined as $D_i = [(i-1)/2k, i/2k)$ except that $D_{2k} = [(2k-1)/2k, 1]$, so the collection $\{D_i\}_{i=1}^{2k}$ forms a partition of the unit interval $[0,1]$. By construction of the densities f_ν, the sign of $f_\nu - 1$ remains constant on each D_i. Let us define (for shorthand) the linear functionals $\varphi_j : L^\infty([0,1]) \to \mathbb{R}$ for each $j \in \{1, \ldots, k\}$ via

$$\varphi_j(\gamma) := \int \gamma(dP_{+j} - dP_{-j}) = \sum_{i=1}^{2k} \int_{D_i} \gamma(x)(f_{+j}(x) - f_{-j}(x))dx = 2 \int_{D_{2j-1} \cup D_{2j}} \gamma(x) g_{\beta,j}(x) dx,$$

where we recall the definitions (61) of the mixture densities $f_{\pm j} = 1 \pm g_{\beta,j}$. Since the set $\mathbb{B}_\infty$ from Theorem 3 is compact, convex, and Hausdorff, the Krein-Milman theorem [42, Proposition 1.2] guarantees that it is equal to the convex hull of its extreme points; moreover, since the functionals $\gamma \mapsto \varphi_j^2(\gamma)$ are convex, the supremum in Theorem 3 must be attained at the extreme points of $\mathbb{B}_\infty([0,1])$. As a consequence, when applying the divergence bound

$$\sum_{j=1}^{k} \left(D_{\rm kl}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\rm kl}\left(M_{-j}^n \| M_{+j}^n\right)\right) \leq 2n(e^\alpha - 1)^2 \sup_{\gamma \in \mathbb{B}_\infty} \sum_{j=1}^{k} \varphi_j^2(\gamma), \tag{64}$$

we can restrict our attention to $\gamma \in \mathbb{B}_\infty$ for which $\gamma(x) \in \{-1, 1\}$.

Now we argue that it is no loss of generality to assume that γ, when restricted to D_i, is a constant (apart from a measure zero set). Fix $i \in [2k]$, and assume for the sake of contradiction that there exist sets $B_i, C_i \subset D_i$ such that $\gamma(B_i) = \{1\}$ and $\gamma(C_i) = \{-1\}$, while $\mu(B_i) > 0$ and $\mu(C_i) > 0$ where μ denotes Lebesgue measure.[4] We will construct vectors γ_1 and $\gamma_2 \in \mathbb{B}_\infty$ and a value $\lambda \in (0,1)$ such that

$$\int_{D_i} \gamma(x) g_{\beta,j}(x) dx = \lambda \int_{D_i} \gamma_1(x) g_{\beta,j}(x) dx + (1-\lambda) \int_{D_i} \gamma_2(x) g_{\beta,j}(x) dx$$

simultaneously for all $j \in [k]$, while on $D_i^c = [0,1] \setminus D_i$, we will have the equivalence

$$\gamma_1|_{D_i^c} \equiv \gamma_2|_{D_i^c} \equiv \gamma|_{D_i^c}.$$

Indeed, set $\gamma_1(D_i) = \{1\}$ and $\gamma_2(D_i) = \{-1\}$, otherwise setting $\gamma_1(x) = \gamma_2(x) = \gamma(x)$ for $x \notin D_i$. For the unique index $j \in [k]$ such that $[(j-1)/k, j/k] \supset D_i$, we define

$$\lambda := \frac{\int_{B_i} g_{\beta,j}(x) dx}{\int_{D_i} g_{\beta,j}(x) dx} \quad \text{so} \quad 1 - \lambda = \frac{\int_{C_i} g_{\beta,j}(x) dx}{\int_{D_i} g_{\beta,j}(x) dx}.$$

By the construction of the function g_β, the functions $g_{\beta,j}$ do not change signs on D_i, and the absolute continuity conditions on g_β specified in equation (58) guarantee $1 > \lambda > 0$, since $\mu(B_i) > 0$ and $\mu(C_i) > 0$. We thus find that for any $j \in [k]$,

$$\begin{aligned}
\int_{D_i} \gamma(x) g_{\beta,j}(x) dx &= \int_{B_i} \gamma_1(x) g_{\beta,j}(x) dx + \int_{C_i} \gamma_2(x) g_{\beta,j}(x) dx \\
&= \int_{B_i} g_{\beta,j}(x) dx - \int_{C_i} g_{\beta,j}(x) dx = \lambda \int_{D_i} g_{\beta,j}(x) dx - (1-\lambda) \int_{D_i} g_{\beta,j}(x) dx \\
&= \lambda \int \gamma_1(x) g_{\beta,j}(x) dx + (1-\lambda) \int \gamma_2(x) g_{\beta,j}(x) dx.
\end{aligned}$$

(Notably, for j such that $g_{\beta,j}$ is identically 0 on D_i, this equality is trivial.) By linearity and the strong convexity of the function $x \mapsto x^2$, then, we find that for sets $E_j := D_{2j-1} \cup D_{2j}$,

$$\begin{aligned}
\sum_{j=1}^{k} \varphi_j^2(\gamma) &= \sum_{j=1}^{k} \left(\int_{E_j} \gamma(x) g_{\beta,j}(x) dx \right)^2 \\
&< \lambda \sum_{j=1}^{k} \left(\int_{E_j} \gamma_1(x) g_{\beta,j}(x) dx \right)^2 + (1-\lambda) \sum_{\nu \in \mathcal{V}} \left(\int_{E_j} \gamma_2(x) g_{\beta,j}(x) dx \right)^2.
\end{aligned}$$

[4] For a function f and set A, the notation $f(A)$ denotes the image $f(A) = \{f(x) \mid x \in A\}$.

Thus one of the densities γ_1 or γ_2 must have a larger objective value than γ. This is our desired contradiction, which shows that (up to measure zero sets) any γ attaining the supremum in the information bound (64) must be constant on each of the D_i.

Having shown that γ is constant on each of the intervals D_i, we conclude that the supremum (64) can be reduced to a finite-dimensional problem over the subset

$$\mathcal{B}_{1,2k} := \left\{ u \in \mathbb{R}^{2k} \mid \|u\|_\infty \leq 1 \right\}$$

of $\mathbb{R}^{2k}$. In terms of this subset, the supremum (64) can be rewritten as the the upper bound

$$\sup_{\gamma \in \mathbb{B}_\infty} \sum_{j=1}^k \varphi_j(\gamma)^2 \leq \sup_{\gamma \in \mathcal{B}_{1,2k}} \sum_{j=1}^k \left(\gamma_{2j-1} \int_{D_{2j-1}} g_{\beta,j}(x)dx + \gamma_{2j} \int_{D_{2j}} g_{\beta,j}(x)dx \right)^2.$$

By construction of the function g_β, we have the equality

$$\int_{D_{2j-1}} g_{\beta,j}(x)dx = -\int_{D_{2j}} g_{\beta,j}(x)dx = \int_0^{\frac{1}{2k}} g_{\beta,1}(x)dx = \int_0^{\frac{1}{2k}} \frac{1}{k^\beta} g_\beta(kx)dx = \frac{c_{1/2}}{k^{\beta+1}}.$$

This implies that

$$\begin{aligned} &\frac{1}{2e^\alpha(e^\alpha-1)^2 n} \sum_{j=1}^k \left(D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right)\right) \leq \sup_{\gamma \in \mathbb{B}_\infty} \sum_{j=1}^k \varphi_j(\gamma)^2 \\ &\leq \sup_{\gamma \in \mathcal{B}_{1,2k}} \sum_{j=1}^k \left(\frac{c_{1/2}}{k^{\beta+1}} \gamma^\top (e_{2j-1} - e_{2j}) \right)^2 = \frac{c_{1/2}^2}{k^{2\beta+2}} \sup_{\gamma \in \mathcal{B}_{1,2k}} \gamma^\top \sum_{j=1}^k (e_{2j-1} - e_{2j})(e_{2j-1} - e_{2j})^\top \gamma, \quad (65) \end{aligned}$$

where $e_j \in \mathbb{R}^{2k}$ denotes the jth standard basis vector. Rewriting this using the Kronecker product $\otimes$, we have

$$\sum_{j=1}^k (e_{2j-1} - e_{2j})(e_{2j-1} - e_{2j})^\top = I_{k\times k} \otimes \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \preceq 2 I_{2k \times 2k}.$$

Combining this bound with our inequality (65), we obtain

$$\sum_{j=1}^k \left(D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right) + D_{\mathrm{kl}}\left(M_{+j}^n \| M_{-j}^n\right)\right) \leq 4n(e^\alpha - 1)^2 \frac{c_{1/2}^2}{k^{2\beta+2}} \sup_{\gamma \in \mathcal{B}_{1,2k}} \|\gamma\|_2^2 = 4c_{1/2}^2 \frac{n(e^\alpha-1)^2}{k^{2\beta+1}}.$$

I Technical arguments

In this appendix, we collect proofs of technical lemmas and results needed for completeness.

I.1 Proof of Lemma 1

Fix an (arbitrary) estimator $\widehat{\theta}$. By assumption (29), we have

$$\Phi(\rho(\theta, \theta(P_\nu))) \geq 2\delta \sum_{j=1}^d \mathbf{1}\{[\mathrm{v}(\theta)]_j \neq \nu_j\}.$$

Taking expectations, we see that

$$\begin{aligned}\sup_{P\in\mathcal{P}} \mathbb{E}_P\left[\Phi(\rho(\widehat{\theta}(Z_1,\ldots,Z_n),\theta(P)))\right] &\geq \max_{\nu\in\mathcal{V}} \mathbb{E}_{P_\nu}\left[\Phi(\rho(\widehat{\theta}(Z_1,\ldots,Z_n),\theta_\nu))\right] \\ &\geq \frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}} \mathbb{E}_{P_\nu}\left[\Phi(\rho(\widehat{\theta}(Z_1,\ldots,Z_n),\theta_\nu))\right] \\ &\geq \frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}} 2\delta \sum_{j=1}^{d} \mathbb{E}_{P_\nu}\left[\mathbf{1}\left\{[\psi(\widehat{\theta})]_j \neq \nu_j\right\}\right],\end{aligned}$$

as the average is smaller than the maximum of a set and using the separation assumption (29). Recalling the definition (30) of the mixtures $P_{\pm j}$, we swap the summation orders to see that

$$\begin{aligned}\frac{1}{|\mathcal{V}|}\sum_{\nu\in\mathcal{V}} P_\nu\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right) &= \frac{1}{|\mathcal{V}|}\sum_{\nu:\nu_j=1} P_\nu\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right) + \frac{1}{|\mathcal{V}|}\sum_{\nu:\nu_j=-1} P_\nu\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right) \\ &= \frac{1}{2}P_{+j}\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right) + \frac{1}{2}P_{-j}\left([\mathsf{v}(\widehat{\theta})]_j \neq \nu_j\right).\end{aligned}$$

This gives the statement claimed in the lemma, while taking an infimum over all testing procedures $\psi : \mathcal{Z}^n \to \{-1,+1\}$ gives the claim (31).

I.2 Proof of unbiasedness for sampling strategy (27)

We compute the expectation of a random variable Z sampled according to the strategy (27); i.e., we compute $\mathbb{E}[Z \mid v]$ for a vector $v \in \mathbb{R}^d$. By scaling, it is no loss of generality to assume that $\|v\|_2 = 1$, and using the rotational symmetry of the ℓ_2-ball, we see it is no loss of generality to assume that $v = e_1$, the first standard basis vector.

Let the function s_d denote the surface area of the sphere in $\mathbb{R}^d$, so that

$$s_d(r) = \frac{d\pi^{d/2}}{\Gamma(d/2+1)} r^{d-1}$$

is the surface area of the sphere of radius r. (We use s_d as a shorthand for $s_d(1)$ when convenient.) Then for a random variable W sampled uniformly from the half of the ℓ_2-ball with first coordinate $W_1 \geq 0$, symmetry implies that by integrating over the radii of the ball,

$$\mathbb{E}[W] = e_1 \frac{2}{s_d}\int_0^1 s_{d-1}\left(\sqrt{1-r^2}\right) r dr.$$

Making the change of variables to spherical coordinates (we use ϕ as the angle), we have

$$\frac{2}{s_d}\int_0^1 s_{d-1}\left(\sqrt{1-r^2}\right) r dr = \frac{2}{s_d}\int_0^{\pi/2} s_{d-1}\left(\cos\phi\right)\sin\phi\, d\phi = \frac{2s_{d-1}}{s_d}\int_0^{\pi/2} \cos^{d-2}(\phi)\sin(\phi)\, d\phi.$$

Noting that $\frac{d}{d\phi}\cos^{d-1}(\phi) = -(d-1)\cos^{d-2}(\phi)\sin(\phi)$, we obtain

$$\int_0^{\pi/2} \cos^{d-2}(\phi)\sin(\phi)\, d\phi = -\frac{\cos^{d-1}(\phi)}{d-1}\Big|_0^{\pi/2} = \frac{1}{d-1}.$$

Thus

$$\mathbb{E}[W] = 2e_1 \frac{(d-1)\pi^{\frac{d-1}{2}}\Gamma(\frac{d}{2}+1)}{d\pi^{\frac{d}{2}}\Gamma(\frac{d-1}{2}+1)} \frac{1}{d-1} = e_1 \underbrace{\frac{2\Gamma(\frac{d}{2}+1)}{\sqrt{\pi}d\Gamma(\frac{d-1}{2}+1)}}_{=:c_d}, \tag{66}$$

where we define the constant c_d to be the final ratio.

Allowing again $\|v\|_2 \le r$, with the expression (66), we see that for our sampling strategy for Z, we have

$$\mathbb{E}[Z \mid v] = v\frac{B}{r}c_d\left(\frac{e^\alpha}{e^\alpha+1} - \frac{1}{e^\alpha+1}\right) = \frac{B}{r}c_d\frac{e^\alpha-1}{e^\alpha+1}.$$

Consequently, the choice

$$B = \frac{e^\alpha+1}{e^\alpha-1}\frac{r}{c_d} = \frac{e^\alpha+1}{e^\alpha-1}\frac{r\sqrt{\pi}d\Gamma(\frac{d-1}{2}+1)}{2\Gamma(\frac{d}{2}+1)}$$

yields $\mathbb{E}[Z \mid v] = v$. Moreover, we have

$$\|Z\|_2 = B \le r\frac{e^\alpha+1}{e^\alpha-1}\frac{3\sqrt{\pi}\sqrt{d}}{4}$$

by Stirling's approximation to the Γ-function. By noting that $(e^\alpha+1)/(e^\alpha-1) \le 3/\alpha$ for $\alpha \le 1$, we see that $\|Z\|_2 \le 4r\sqrt{d}/\alpha$.

I.3 Proof of unbiasedness for sampling strategy (28)

We compute conditional expectations for each of the uniform quantities in the sampling scheme (28). This argument is based on Corollary 3.7 of Duchi et al. [15], but we present it here for clarity and completeness. In each summation to follow, we implicitly assume that $z \in \{-1,1\}^d$ and that $x \in \{-1,1\}^d$, as the general result follows by scaling of these quantities. We consider first the case that d is odd. In this case, we have by symmetry that

$$\begin{aligned}\sum_{z:\langle z,x\rangle \ge 0} z &= \sum_{z:\langle z,x\rangle=1} z + \sum_{z:\langle z,x\rangle=3} z + \cdots + \sum_{z:\langle z,x\rangle=d} z \\ &= \left[\binom{d-1}{\frac{d-1}{2}} - \binom{d-1}{\frac{d+1}{2}}\right]x + \left[\binom{d-1}{\frac{d+1}{2}} - \binom{d-1}{\frac{d+3}{2}}\right]x + \cdots + \binom{d-1}{d-1}x = \binom{d-1}{\frac{d-1}{2}}x,\end{aligned}$$

and as $|\{z \in \{-1,1\}^d \mid \langle z,x\rangle > 0\}| = 2^{d-1}$, we obtain that for $Z \sim \text{Uniform}(\{-1,1\}^d)$ we have

$$\mathbb{E}[Z \mid \langle Z,x\rangle \ge 0] = \frac{1}{2^{d-1}}\binom{d-1}{\frac{d-1}{2}}x \quad \text{and} \quad \mathbb{E}[Z \mid \langle Z,x\rangle \le 0] = -\frac{1}{2^{d-1}}\binom{d-1}{\frac{d-1}{2}}x.$$

Similarly, for d even, we find that

$$\begin{aligned}\sum_{z:\langle z,x\rangle \ge 0} z &= \sum_{z:\langle z,x\rangle=2} z + \sum_{z:\langle z,x\rangle=4} z + \cdots + \sum_{z:\langle z,x\rangle=d} z \\ &= \left[\binom{d-1}{\frac{d}{2}} - \binom{d-1}{\frac{d}{2}+1}\right]x + \left[\binom{d-1}{\frac{d}{2}+1} - \binom{d-1}{\frac{d}{2}+2}\right]x + \cdots + \binom{d-1}{d-1}x = \binom{d-1}{\frac{d}{2}}x.\end{aligned}$$

Noting that the set $\{z \in \{-1,1\}^d \mid \langle z, x\rangle \geq 0\}$ has cardinality $2^{d-1} + \frac{1}{2}\binom{d}{d/2}$ (because we must consider vectors z such that $\langle z, x\rangle = 0$), we find that for d even we have

$$\mathbb{E}[Z \mid Z^\top x \geq 0] = \frac{1}{2^{d-1} + \frac{1}{2}\binom{d}{d/2}} \binom{d-1}{\frac{d}{2}} x \quad \text{and} \quad \mathbb{E}[Z \mid Z^\top x \leq 0] = -\frac{1}{2^{d-1} + \frac{1}{2}\binom{d}{d/2}} \binom{d-1}{\frac{d}{2}} x.$$

Inverting the constant multipliers on the vectors x in the preceding equations shows that the strategy (28) is unbiased.